



Ivanti Connect Secure Administration Guide
9.1R15

Copyright Notice

This document is provided strictly as a guide. No guarantees can be provided or expected. This document contains the confidential information and/or proprietary property of Ivanti, Inc. and its affiliates (referred to collectively as "Ivanti") and may not be disclosed or copied without prior written consent of Ivanti.

Ivanti retains the right to make changes to this document or related product specifications and descriptions, at any time, without notice. Ivanti makes no warranty for the use of this document and assumes no responsibility for any errors that can appear in the document nor does it make a commitment to update the information contained herein. For the most current product information, please visit www.lvanti.com.

Copyright © 2022, Ivanti, Inc. All rights reserved.

Protected by patents, see <https://www.ivanti.com/patents>.

Table of Contents

Revision History	11
Preface	18
Document conventions	18
Requesting Technical Support	19
Introduction	21
About the Ivanti Connect Secure Administration Guide	21
Scope	21
Ivanti Connect Secure Documentation and Resources	21
Key Terms and Concepts	21
Ivanti Connect Secure Overview	23
Using Ivanti Connect Secure for Securing Traffic	24
Authenticating Users with Existing Servers	26
Using Host Checker to Protect from Threats	28
Configuring Ivanti Connect Secure	30
Introducing the Pulse Clients	31
User Verification and Key Concepts	33
Verifying User Accessibility	33
General Access Management	43
Access Management Overview	43
Policies, Rules & Restrictions, and Conditions Overview	43
Policies, Rules & Restrictions, and Conditions Evaluation	45
Dynamic Policy Evaluation	48
Specifying Source IP Access Restrictions	50
Specifying Browser Access Restrictions	54
Specifying Certificate Access Restrictions	57
Specifying Password Access Restrictions	59
Specifying Session Limits	60
IF-MAP Federation Overview	63
IF-MAP Federation Details	65
Task Summary: Configuring IF-MAP Federation	68
Configuring IF-MAP Server Settings	69
Configuring the IF-MAP Federation Client	69
IF-MAP Federated Network Timing Considerations	70
Session-Export and Session-Import Policies	70
Configuring Session-Export Policies	73
Session-Import Policies	76
Troubleshooting the IF-MAP Federated Network	76
Viewing Active Users on the IF-MAP Client	77
Trusted Server List	77
User Roles	80
User Roles Overview	80
Configuring General Role Options	84
Role Restrictions	85
Specifying Role-Based Source IP Aliases	86

Specifying Role Session Options	86
Customizing the Welcome Page	91
Optimized Interface for the Apple iPad	96
Defining Default Options for User Roles	98
Customizing Messages	100
Customizing UI Views for User Roles	100
Virtual Desktop Resource Profiles	104
Virtual Desktop Resource Profile Overview	104
Configuring a Citrix XenDesktop Resource Policy	104
Configuring a VMware View Manager Resource Profile	106
Defining Bookmarks for a Virtual Desktop Profile	107
Configuring the Client Delivery	108
Connecting to the Servers	109
Authentication and Directory Servers	110
AAA Server Overview	110
AAA Traffic Management	111
Using the Local Authentication Server	115
Using Active Directory	121
JITC AAA Certification	130
Understanding Multidomain User Authentication	131
Understanding Active Directory and Windows NT Group Information Support	133
Join Domain for Active Directory-based Authentication Server Without Using a Domain Admin Account	134
Using the Certificate Server	134
Using an LDAP Server	136
Using the LDAP Password Management Feature	141
Using an MDM Server	146
Using a RADIUS Server	149
Using an ACE Server	168
Using the SAML Server	172
Authentication Realms	193
Understanding Authentication Realms	193
Creating an Authentication Realm	193
Role Mapping Rules	195
Specifying Role Mapping Rules for an Authentication Realm	196
Machine Authentication for Pulse Client Connections	198
Connection Realm and Role Preferences for Machine Authentication	199
Configuring Role Mapping Rules based on Geo Location Custom Expressions	201
Using the LDAP Server Catalog	205
Customizing User Realm UI Views	210
Single Sign-On	212
About Single Sign-On	212
About Multiple Sign-In Credentials	213
Task Summary: Configuring Multiple Authentication Servers	213
Task Summary: Enabling SSO to Resources Protected by Basic Authentication	214
Enabling SSO to Resources Protected by NTLM	214

Multiple Sign-In Credentials Execution	215
Adaptive Authentication	221
Overview	221
Adaptive Authentication User Flow	221
Dashboard and Reports	224
Troubleshooting	225
Synchronizing User Records	227
About User Record Synchronization	227
Enabling User Record Synchronization	229
Configuring the User Record Synchronization Authentication Server	229
Configuring the User Record Synchronization Server	230
Configuring the User Record Synchronization Client	231
Configuring the User Record Synchronization Database	231
Scheduling User Record Synchronization Backup	233
Host Checker	235
Sign-In Policies	236
About Sign-In Policies	236
Task Summary: Configuring Sign In Pages	237
About Configuring Sign In Policies	238
Configuring User Sign In Policies	238
Configuring Fallback Authentication Server	241
About Sign-In Notifications	241
Configuring and Implementing Sign-in Notifications	242
Defining Authorization-Only Access Policies	245
Configuring Sign-In Pages	247
Resource Profiles	250
Resource Profiles	250
Resource Profile Components	250
Defining Resource Profile Resources	253
Defining Resource Profile Autopolicies	254
Defining Resource Profile Roles	255
Defining Resource Profile Bookmarks	256
Resource Profile Templates	257
SAML Single Sign-on	258
Ivanti Connect Secure SAML 2.0 SSO Solutions	258
SAML 2.0 Configuration Tasks	270
Example: Implementing SAML 2.0 Web Browser SSO for Google Apps	304
Using SAML AuthnContext Class Variables in Role Mapping and Web ACL Rules	313
Investigating a "No valid assertion found in SAML response" Error	321
Ivanti Connect Secure SAML 1.1 Support	323
Device Access Management Framework	351
Hosted Java Applets Templates	353
About Hosted Java Applet Templates	353
Task Summary: Hosting Java Applets	353
Uploading Java Applets to Ivanti Connect Secure	354
Signing Uploaded Java Applets	355

Creating HTML Pages That Reference Uploaded Java Applets	355
Accessing Java Applet Bookmarks	355
Creating a Hosted Java Applet Resource Profile	356
Configuring Hosted Java Applet Resource Profile Bookmarks	358
Creating Hosted Java Applets Bookmarks Through the User Roles Page	360
Required Attributes for Uploaded Java Applets	361
Required Parameters for Uploaded Java Applets	362
Resource Policies	364
Resource Policies	364
Resource Policy Components	365
Specifying Resources for a Resource Policy	365
Resource Policy Evaluation	368
Creating Detailed Rules for Resource Policies	370
Writing a Detailed Rule for Resource Policies	371
Customizing Resource Policy UI Views	374
Citrix Templates	375
About Citrix Templates	375
Comparing Access Mechanisms for Configuring Citrix	375
Creating Resource Profiles for Citrix Storefront Server	378
Lotus iNotes Templates	381
Creating Resource Profiles Using the Lotus iNotes Template	381
Microsoft OWA Templates	384
Creating Resource Profiles Using the Microsoft OWA Template	384
Microsoft RDWeb HTML5 Templates	387
Creating Resource Profiles Using the Microsoft RDWeb Template	387
Microsoft Sharepoint Templates	390
Creating Resource Profiles Using the Microsoft Sharepoint Template	390
Web Rewriting	393
File Rewriting	394
PleasSecure Application Manager	395
Secure Application Manager Overview	395
Task Summary: Configuring PSAM	395
PSAM Recommended Operation	397
Debugging PSAM Issues	398
About PSAM Resource Profiles	398
Creating PSAM Client Application Resource Profiles	399
Creating PSAM Destination Network Resource Profiles	402
Specifying Applications and Servers for PSAM to Secure	403
Specifying Applications that Need to Bypass PSAM	405
Specifying Role-Level PSAM Options	406
Specifying Application Servers that Users Can Access	408
Specifying Resource Level PSAM Options	410
JSAM Overview	410
Task Summary: Configuring JSAM	411
Using JSAM for Client/Server Communications	412
Configuring a PC that Connects Through a Proxy Web Server	417

Determining the Assigned Loopback Address	417
Configuring External DNS Servers and User Machines	419
JSAM Linux and Macintosh Support	420
Standard Application Support: MS Outlook	420
Standard Application Support: Lotus Notes	422
Configuring the Lotus Notes Client	423
Enabling Citrix Published Applications on the Citrix Native Client	424
Enabling Citrix Secure Gateways	426
Creating a JSAM Application Resource Profile	427
Specifying Applications for JSAM to Secure	431
Specifying Role Level JSAM Options	434
Automatically Launching JSAM	435
Specifying Application Servers that Users Can Access	437
Specifying Resource Level JSAM Options	437
Terminal Services	439
About Terminal Services	439
Task Summary: Configuring the Terminal Services Feature	440
Terminal Services Execution	442
Configuring Citrix to Support ICA Load Balancing	444
About Terminal Services Resource Profiles	447
Configuring a Windows Terminal Services Resource Profile	448
Defining a Hosted Java Applet Autopolicy	450
Defining a Bookmark for a Windows Terminal Services Profile	453
Creating a Windows Terminal Services Bookmark Through the User Roles Page	455
Defining Display Options for the Windows Terminal Services Session	456
Defining SSO Options for the Windows Terminal Services Session	457
Defining Application Settings for the Windows Terminal Services Session	457
Defining Device Connections for the Windows Terminal Services Session	459
Defining Desktop Settings for the Windows Terminal Services Session	461
Creating a Citrix Terminal Services Resource Profile Using Default ICA Settings	462
Defining a Bookmark for a Citrix Profile Using Default ICA Settings	463
Creating a Citrix Terminal Services Bookmark Through the User Roles Page	465
Defining Display Options for the Citrix Terminal Services Session	466
Defining SSO Options for the Citrix Terminal Services Session	467
Defining Application, Auto-Launch, and Session Reliability Settings for the Citrix Terminal Services Session	468
Defining Device Connections for the Citrix Terminal Services Session	470
Creating a Citrix Resource Profile That Uses a Custom ICA File	471
Defining a Bookmark for a Citrix Profile Using a Custom ICA File	474
Creating a Citrix Profile That Lists Published Applications	475
Defining a Bookmark for a Citrix Profile Listing Applications	477
Creating Session Bookmarks to Your Terminal Server	480
Creating Advanced Terminal Services Session Bookmarks	481
Defining Screen Size and Color Depth Options for the Terminal Services Session	483
Defining SSO Options for the Terminal Services Session	484
Defining Application Settings for the Terminal Services Session	486

Defining Device Connections for the Terminal Services Session	487
Defining Desktop Settings for the Terminal Services Session	489
Creating Links from an External Site to a Terminal Services Session Bookmark	490
Specifying General Terminal Services Options	498
Configuring Terminal Services Resource Policies	502
Specifying the Terminal Services Resource Option	504
Using the Remote Desktop Launcher	505
Remote Desktop and Telnet/SSH via HTML5 Access	507
Configuring the HTML5 Access Feature	507
VPN Tunneling	526
Enterprise Onboarding	527
Configuring Enterprise Onboarding	527
Managing Onboarded Devices	548
Cloud Secure	549
Network and Host Administration	550
Network and Host Administration Overview	550
Configuring the Internal Port	550
Configuring the External Port	554
Using the Internal and External Ports	557
Using the Management Port	558
Configuring VLAN Ports	564
Using Virtual Ports	567
Configuring the System Date and Time	571
Configuring Network Services	574
Configuring NTP and Other Services Traffic Over Any Physical Interface	577
Managing the Routes Table	578
Managing the Hosts Table	579
Proxy Server Configuration	579
Managing the ARP Table	581
Managing the Neighbor Discovery Table	582
Using IPv6	583
Configuring SSL Options	596
Enabling Granular Cipher Selection for Setting the Security Options	596
Configuring Health Check Options	608
Configuring Miscellaneous Security Options	609
Configuring Custom HTTP Headers	615
Configuring NCP and JCP	617
Using the User Record Synchronization Feature	618
Using IKEv2 Security	625
Using the Mobile Options	649
Using the Advanced Client Configuration Feature	650
Using the Traffic Segregation Feature	651
Using the Serial Port	654
Certificate Security Administration	659
Understanding Digital Certificate Security	659
Using Device Certificates	660

Using Trusted Client CAs	668
Using Trusted Server CAs	679
Using Code-Signing CAs	681
Using Client Auth Certificates	685
Mapping Resource Policies to the Certificate	689
Mapping a Client Authentication Auto-Policy	691
Checking Certificate Expiry	691
Elliptic Curve Cryptography	694
Understanding ECC Certificates	694
Example: Assigning an ECC P-256 Certificate to an External Virtual Port and Giving Preference to Suite B Ciphers	695
Configuration File Administration	713
Configuration File Administration Overview	713
Configuring Archiving for System Logs, Configuration Files, and Snapshots	714
Using the Configuration Backup and Restore Feature	718
Using the Import/Export Feature for Binary System Configuration Files	720
Using the Import/Export Feature for Binary User Configuration Files	725
Using the Import/Export Feature for XML Configuration Files	727
Example: Using the Configuration XML File Import/Export Feature to Add Multiple Users	734
Guidelines for Modifying Configuration XML Files	736
Using the Push Configuration Feature	745
System Maintenance	758
Using the System Maintenance Pages	758
Configuring System Maintenance Options	758
Upgrading the System Software	762
Downloading Client Installer Files	766
Restarting, Rebooting, and Shutting Down the System	767
Testing Network Connectivity	769
Logging and Monitoring	770
Logging Overview	770
Configuring Events to Log	771
Enabling Client-Side Logging	773
Enabling and Viewing Client-Side Log Uploads	775
Configuring SNMP	777
Configuring Syslog	789
Configuring Advanced Settings	791
Displaying System Status	793
Displaying Hardware Status	797
Using Software RAID PSA7000	798
LCD Display	800
Displaying Active Users	804
Displaying System Logs	806
Using Log Filters	809
Displaying User Access Statistics	813
Troubleshooting Tools	815
Using the Admin Console Troubleshooting Tools	815

Clustering	845
Delegating Administrator Roles	846
About Delegating Administrator Roles	846
Creating and Configuring Administrator Roles	847
Specifying Management Tasks to Delegate	848
Deployments with IDP	851
About IDP	851
IDP Deployment Scenarios	852
Configuring Ivanti Connect Secure to Interoperate with IDP	852
Identifying and Managing Quarantined Users Manually	853
Dashboard and Reports	855
Pulse One Integration	856
Overview	856
Pulse Workspace Handlers	860
Customizable Admin and End-User UIs	863
Customizable Admin and End-User UIs	863
Customizable End-User Interface Elements Overview	864
REST Support for Ivanti Connect Secure	864
FIPS Level 1 Support (Software FIPS)	875
Understanding Ivanti FIPS Level 1 Support	875
Enabling FIPS Level 1 Support	876
Turning Off FIPS Level 1 Support from the Serial Console	878
Installing a Self-Signed Certificate from the Serial Console	879
Supported Cipher Suites when FIPS Level 1 Support is Enabled and Disabled	881
Supported Cipher Suites When FIPS Level 1 Support is Enabled	886
Compression	889
About Compression	889
Enabling System-Level Compression	891
Localization	892
About Multi-Language Support for Ivanti Connect Secure	892
Encoding Files for Multi-Language Support	892
Localizing the User Interface	893
Localizing Custom Sign-In and System Pages	893
Smart Phones	894
Smart Phones	894
Configuring Connect Secure for PDAs and Handhelds	894
Defining Client Types	896
Enabling ActiveSync for Handheld Devices	899
Custom Expressions and System Variables	902
Using Custom Expressions in Rule Configuration	902

Revision History

The following table lists changes to this document from the previous release:

Feature	Add/Update	Effective Release	Notes
HTML 5 - End User bookmarks and Session Recording	Updated Section : Remote Desktop and Telnet/SSH via HTML5 Access	9.1R15	NA
Deprecation of features	<p>Removed content for following:</p> <ul style="list-style-type: none"> • Auth Servers: Siteminder, NIS, Anonymous, Novell eDirectory, and iPlanet • Host Checker: Support for Solaris, Cache cleaner, and Statement of Health • Re-writer: Citrix Web Interface/JICA, Citrix storefront below 3.1, Microsoft OWA 2000, 2003, 2007, Lotus inotes 5, 6, 6.5 • Sensors 	9.1R15	Upgrade to this version may fail if the configurations are not removed. Refer KB45044.

Feature	Add/Update	Effective Release	Notes
	<ul style="list-style-type: none"> • Cloud Application visibility • Telnet/SSH Resource profile • UNIX(NFS) file share • Software defined perimeter (SDP) • Pulse Collaboration 		
Rebranding	Rebranding the entire guide text to make it Ivanti, Ivanti Policy Secure and Ivanti Connect Secure specific as applicable	9.1R14	
Integrity Checker Tool	Configuring Miscellaneous Security Options	9.1R12	
Advanced HTML5 Enhancements	Defining Bookmarks for HTML5 Access	9.1R12	
Intune integration enhancements	Resource Profile	9.1R12	
Syslog Server enhancement	Configuring Syslog	9.1R12	
Enable Pulse Client Components removal Tool option	Updated the Configuring System Maintenance Options section.	9.1R11.5	

Feature	Add/Update	Effective Release	Notes
Advanced HTML5, Auto system reboot	Updated sections Configuring System Maintenance Options , Remote Desktop and Telnet/SSH via HTML5 Access	9.1R11	
Host header validation	Updated the Configuring Miscellaneous Security Options section.	9.1R10	
Advanced HTML5 feature	Updated the Displaying System Status section.	9.1R10	
Multiple users for SNMP	Updated the Configuring SNMP topic.	9.1R9	
SNMP Diagnostics log	Added the Using the SNMP Diagnostic Log topic.	9.1R9	
ESP Tunnel for mixed mode	Added the Defining ESP Tunnel for Mixed Mode Traffic topic.	9.1R9	
Advanced HTML5 feature	Updated the Remote Desktop and Telnet/SSH via HTML5 Access chapter.	9.1R9	
Remote microphone support in WTS	Updated the Terminal Services chapter.	9.1R9	
Improvement	Updated the "Credential Provider Authentication for Connect Secure" topic in VPN Tunneling .	9.1R8	
Improvement	Updated the Default Settings for Administrators topic.	9.1R8	

Feature	Add/Update	Effective Release	Notes
Host Checker Policies and Supported Platforms	Added the sections "Policies" and "Supported Platform Matrix" in Host Checker .	9.1R8	
Overload protection	Updated the sections Configuring System Maintenance Options and Configuring Events to Log	9.1R8	
Support NTP pool of servers	Updated the section Configuring the System Date and Time	9.1R8	
Show users by access type	Updated the Displaying System Status section.	9.1R8	
Route Table controls	Updated the Managing the Routes Table section	9.1R8	
REST Monitor	Added the Using the REST Monitor section.	9.1R8	
UEBA package for new installation	Updated the Configuring ICS for Enabling Behavioral Analytics and System Variables sections.	9.1R8	
Dashboard graph	Updated the "Dashboard and Report Overview" section in Dashboard and Reports .	9.1R8	
Enable/Disable FQDN ACL	Updated the "Defining VPN Tunneling Access Control Policies" section in VPN Tunneling .	9.1R7	
Monitoring HTML5 Sessions	Added the Monitoring HTML5 Sessions section.	9.1R7	

Feature	Add/Update	Effective Release	Notes
Password Management for OpenLDAP	Updated the LDAP Password Management Feature Overview section.	9.1R5	
Microsoft Intune Integration	Updated the Using an MDM Server section.	9.1R5	
Enhancements to Health Check for load-balancing	Updated the "Health Checking a Server from a Load Balancer" section in Clustering .	9.1R4	
Document improvement	Updated the "WAN Clustering" section in Clustering , "Using Endpoint Security Assessment Plug-In" section in Host Checker , and IPv6 Support Overview section.	9.1R3, ver 1.1	
Connect to nearest available DC	Updated the Configuring Authentication with an LDAP Server section.	9.1R3	
Consolidated system and troubleshooting logs	Added the Using Log Selection section.	9.1R3	
Control copy/paste option for a user from a HTML5 session	Updated the Using the Debug Log and Creating HTML5 Enduser Bookmarks sections.	9.1R3	
Troubleshooting enhancements	Updated the Using the Debug Log section and "Configuring Cluster Node Monitoring" section in Clustering .	9.1R3	

Feature	Add/Update	Effective Release	Notes
Enhancements to Local Authentication Server default password	Updated the Configuring the Local Authentication Server section.	9.1R3	
Restricting access to default resource policies	Updated the Resource Policies and Defining a Resource Profile sections.	9.1R3	
Setting MSS value for TCP connections	Updated the Configuring Network Services section.	9.1R3	
Network Connect Client	Updated the "Defining VPN Tunneling Role Settings" section in VPN Tunneling .	9.1R3	
SP-Initiated SAML SSO	Added the Configuring Service Provider Initiated SAML SSO section and updated the Configuring a SAML SSO Resource Policy for Gateway Mode Deployments section.	9.1R2	
MS RDWeb HTML5 Access	Added the Microsoft RDWeb HTML5 Templates chapter.	9.1R2	
Backup configs and archived logs on AWS S3/Azure Storage	Updated the Configuring Archiving for System Logs, Configuration Files, and Snapshots section.	9.1R2	
OPSWAT SDK V3 to V4 Migration	Updated the "OPSWAT SDK V3 to V4 Migration" section in Host Checker .	9.1R2	
Steelhead Mobile Controller options		8.0R6	

Feature	Add/Update	Effective Release	Notes
AD Configuration (Legacy Mode)		9.1R1	
DNS traffic on any physical interface	Updated the Configuring Network Services section.	9.1R1	
Support for client-name parameter in HTML5 Access	Updated the Launching Custom Page via HTML5 Access section.	9.1R1	
Authentication failure management	Updated the Using the Local Authentication Server section.	9.1R1	
TOTP enhancement	Updated the Configuring the TOTP Authentication Server Settings section.	9.1R1	
Software Defined Perimeter	Added Introduction section	9.1R1	

Preface

Document conventions

The document conventions describe text formatting conventions, command syntax conventions, and important notice formats used in Ivanti technical documentation.

Text formatting conventions

Text formatting conventions such as boldface, italic, or Courier font may be used in the flow of the text to highlight specific words or phrases.

Format	Description
bold text	Identifies command names
	Identifies keywords and operands
	Identifies the names of user-manipulated GUI elements
	Identifies text to enter at the GUI
<i>italic text</i>	Identifies emphasis
	Identifies variables
	Identifies document titles
Courier Font	Identifies command output
	Identifies command syntax examples

Command syntax conventions

Bold and italic text identify command syntax components. Delimiters and operators define groupings of parameters and their logical relationships.

Convention	Description
bold text	Identifies command names, keywords, and command options.
<i>italic text</i>	Identifies a variable.
[]	Syntax components displayed within square brackets are optional.

Convention	Description
	Default responses to system prompts are enclosed in square brackets.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	A vertical bar separates mutually exclusive elements.
< >	Non-printing characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, member[member...].
\	Indicates a “soft” line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.
bold text	Identifies command names, keywords, and command options.

Notes and Warnings

Note, Attention, and Caution statements might be used in this document.



A Note provides a tip, guidance, or advice, emphasizes important information, or provides a reference to related information.

An Attention statement indicates a stronger note, for example, to alert you when traffic might be interrupted or the device might reboot.

A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.

Requesting Technical Support

Technical product support is available through the Ivanti Global Support Center. If you have a support contract, file a ticket with the Ivanti Global Support Center.

- Product warranties—For product warranty information, visit <https://support.pulsesecure.net/product-service-policies/>

Self-Help Online Tools and Resources

For quick and easy problem resolution, Ivanti provides an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://support.pulsesecure.net>
- Search for known bugs: <https://support.pulsesecure.net>
- Find product documentation: <https://www.ivanti.com/support/product-documentation>
- Download the latest versions of software: <https://support.pulsesecure.net>
- Open a case online in the CSC Case Management tool: <https://support.pulsesecure.net>
- To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool:
<https://support.pulsesecure.net>

For important product notices, technical articles, and to ask advice:

- Search the Knowledge Center for technical bulletins and security advisories:
<https://kb.pulsesecure.net>
- Ask questions and find solutions at the Pulse Community online forum:
<https://community.pulsesecure.net>

Opening a Case with Ivanti Global Support Center

You can open a case with Ivanti Global Support Center on the Web or by telephone.

- Use the Case Management tool in the Ivanti Global Support Center at
<https://support.pulsesecure.net>.
- Call 1-844 751 7629 (Toll Free, US).
- For international or direct-dial options in countries without toll-free numbers, see
<https://support.pulsesecure.net/support/support-contacts/>

Introduction

About the Ivanti Connect Secure Administration Guide

This guide is designed for network administrators to configure and maintain a Ivanti Connect Secure device. To use this guide, you need a broad understanding of networks in general and the Internet in particular, networking principles, and network configuration.

Scope

The Ivanti Connect Secure Administrator Guide provides detailed information on configuring, authenticating, securing, managing, and troubleshooting Ivanti Connect Secure and Pulse Client in your environment. Before you configure your environment, it is mandatory to walk through the *Supported Platforms Guide* and the *License Management Guide*.

Ivanti Connect Secure Documentation and Resources

The Ivanti Connect Secure documentation set includes multiple separate deliverables in the web HTML format. The publications are available at <https://www.ivanti.com/support/product-documentation>.

Key Terms and Concepts

Glossary Acronyms	Description
AAA Server	AAA is expanded as Authentication, Authorization, and Accounting. AAA Server is a server program which provides any of the AAA services, viz, Authentication, Authorization or Accounting.
Access	Refers to the level and the extent of a service's functionality or data that a user is entitled to use.
CIE	Content Intermediation Engine. An advanced parser and rewriter that retrieves Web-based content from internal Web servers and changes URL references and Java socket calls.
Cipher	Cipher is an algorithm for performing encryption or decryption. It is a series of well-defined steps that can be followed as a procedure.

Glossary Acronyms	Description
Compression	A method that is followed by the ICS to improve the performance by compressing common types of Web and file data such as HTML files, Word documents, and images
Digital Certificates	Digital Certificates are issued by Certificate Authority (CA). A digital certificate validates the ownership of a public key with subject name in the certificate.
DMI	Device Management Interface is an XML-RPC based protocol that is used to manage Ivanti Connect Secure devices.
HMAC Key	Hash Message Authentication Code is a specific type of message authentication code (MAC) hashed to identify individual devices to the application.
Host Checker	Host Checker is an endpoint security-based feature, which performs security and system integrity checks that pre-qualify endpoints before allowing access to the network's resources.
IDPS	Intrusion detection and prevention sensor monitors networks to detect suspicious and anomalous network traffic based on specific rules defined in IDP rule bases.
IF-MAP	Interface for Metadata Access Point is a protocol defined by the Trusted Network Connect Working Group (TNC-WG) as a standard interface between different network elements and devices.
Localization	The multi-language support provided by the ICS for file encoding, end-user interface display, and customized sign-in and system pages
Non-broadcast SSID	Non-broadcast Service Set Identifier prevents unauthorized users from being able to detect the wireless network from their wireless clients.
Realm	Specifies the authentication and authorization mechanisms (including Host Checker policies) associated with a given sign-in URL.
Roles	Specifies the user privileges and access mechanism based on the information returned by the realm's directory or the user's name.
Sign-In Policy	Sign-in policies define the URLs that users and administrators use to access the device and the sign-in pages.
SMS	System Management Server provides automatic updates to non-compliant software.

Ivanti Connect Secure Overview

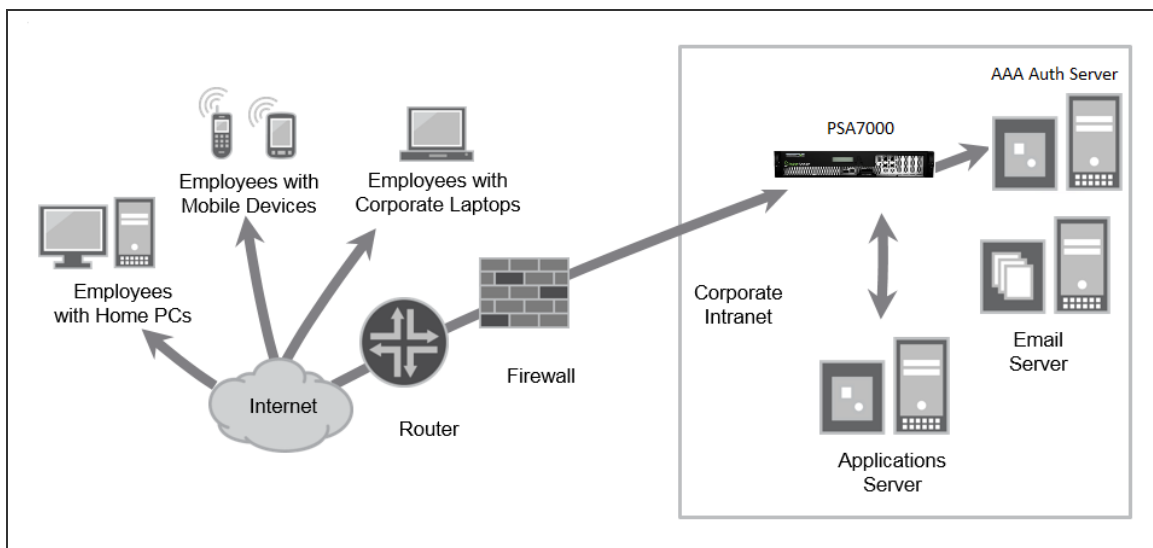
Ivanti Connect Secure gives employees, partners and customers secure and controlled access to corporate data and applications. The applications include file servers, web servers, native messaging, and hosted servers outside your trusted network.

The organization home page can be accessed by employees, partners and customers through a web browser with SSL support and an Internet connection. The page allows the users to:

- Securely browse web or file servers
- Use HTML-enabled enterprise applications
- Start the client/server application proxy
- Begin a Windows or Citrix terminal session
- Access corporate e-mail servers
- Start a secured Layer 3 tunnel

How Ivanti Connect Secure Works

Ivanti Connect Secure authorizes the resources that are accessed by users through an extranet session hosted by the appliance. Ivanti Connect Secure intermediates the data that flows between external users and the company's internal resources to provide robust security. The following diagram is an example of Ivanti Connect Secure within a LAN environment.



During the process of intermediation, the ICS receives secure requests from the external, authenticated users and makes the request to the internal resources on behalf of the users. By intermediating, the need to deploy extranet toolkits in traditional demilitarized zones (DMZ) or provision a remote access VPN for employees is eliminated.

Ivanti Connect Secure Benefits

Ivanti Connect Secure offers high standard configurable solutions. Ivanti Connect Secure:

- Intermediates access to multiple types of applications and resources. These include web-based enterprise applications, Java applications, file shares, terminal hosts, and other client/server applications such as Microsoft Outlook, Lotus Notes, the Citrix XenApp and Smart Phones. Additionally, administrators can provision an access method that allows full Layer 3 connectivity, which provides the same level of access that a user would get if they were on the corporate LAN.
- Fine tunes the user access to the appliance, resource types, or individual resources based on factors such as group membership, source IP address, certificate attributes, and endpoint security status. For example, you can use the dual-factor authentication and client-side digital certificates to authenticate users and use LDAP group membership to authorize users to access individual applications.
- Assesses the security status of your users' computers by checking for endpoint defense tools such as current antivirus software, firewalls, and security patches. You can then allow or deny users access to the appliance, resource types, or individual resources based on the computer's security status.
- Acts as a secure application Layer gateway intermediating all requests between the public Internet and internal corporate resources. All requests that enter the ICS are encrypted by the end user's browser using SSL/TLS. Because the ICS provides a robust security layer between the public Internet and internal resources, administrators do not need to constantly manage security policies and patch security vulnerabilities for numerous different applications and web servers deployed in the public-facing DMZ.

Using Ivanti Connect Secure for Securing Traffic

Ivanti Connect Secure provides secure access to different types of applications, servers and other resources through its remote access mechanism. Simply select both the resources you want to secure and the appropriate access mechanism.

As an example, if you want secure access to Microsoft Outlook, you can use the Secure Application Manager (SAM). The Secure Application Manager intermediates traffic to the client/server applications including Microsoft Outlook, Lotus Notes, and Citrix. Otherwise, if you want to secure access to your company Intranet websites, you can use the web rewriting feature. This feature uses the ICS's Content Intermediation Engine to intermediate traffic to web-based applications and web pages.

Intermediating Traffic Types

The remote access mechanism is integrated with the ICS to intermediate the following types of traffic, the application and the resources that it handles.

- **Web applications and web pages:** Use the web rewriting feature to intermediate web page type of content. The web rewriting feature includes templates that enables you to easily configure access to applications such as Citrix, OWA, Lotus iNotes, and SharePoint. In addition, you can use the web rewriting custom configuration option to intermediate traffic from a wide variety of additional web-based applications, web pages, and custom-built web applications.
- **Web applications using Java applets:** Use the hosted Java applets feature to intermediate this type of content. This feature enables the user to host Java applets and the HTML pages that they reference directly on Ivanti Connect Secure rather than maintaining a separate Java server.
- **File servers and directories using file traffic:** Uses the file rewriting feature to intermediate and dynamically "webify" access to file shares. The file rewriting feature enables you to secure traffic to a variety of Windows and UNIX based servers, directories, and file shares.
- **Client/server applications:** Use the Secure Application Manager (SAM) feature to intermediate this type of content. SAM comes in two varieties (PSAM and JSAM). The PSAM and JSAM features include templates that enable you to easily configure access to applications such as Lotus Notes, Microsoft Outlook, NetBIOS file browsing, and Citrix. In addition, you can use the PSAM and JSAM custom configuration options to intermediate traffic from a wide variety of additional client/server applications and destination networks.
- **Windows Terminal Servers and Citrix server terminal emulation sessions:** Use the Terminal Services feature to intermediate this type of content. This feature enables you to easily configure access to Windows Terminal Servers, Citrix XenApp and StoreFront servers. You can also use this feature to deliver the terminal services clients directly from the ICS, eliminating the need to use another web server to host the clients.

- **All network traffic:** Use the VPN Tunneling feature to create a secure, Layer 3 tunnel over the SSL connection, allowing access to any type of application available on the corporate network. This feature easily connects remote users into your network by tunneling network traffic over port 443, enabling the users with complete access to all network resources without configuring access to individual servers, applications, and resources. Layer 3 VPN tunnels can be initiated using the integrated Network Connect client and the Pulse desktop and mobile clients.

Authenticating Users with Existing Servers

You can easily configure Ivanti Connect Secure to use your company's existing servers to authenticate your end users. Users need not create a new username and password to access the device.

The ICS supports integration with LDAP, RADIUS, Windows NT Domain, Active Directory, SAML, and RSA ACE/Servers.

Alternatively, if you do not want to use one of these standard servers, you can store usernames and credentials directly on the ICS and use it as an authentication server. In addition, you can choose to authenticate users based on attributes contained in authentication assertions generated by SAML authorities or client-side certificates.



Pulse Mobile client supports only one case of dual-factor authentication, in which the client certificate is the primary, while the local authorization is the secondary.

Using Client-side Authorization to Control Access

In addition to using authentication servers to control access to Ivanti Connect Secure, you can control access to the ICS and the resources it intermediates using a variety of additional client-side checks. Ivanti Connect Secure enables you to create a multi-layered approach to protect itself and your resources by doing the following:

1. As a first step, perform pre-authentication checks that control user access to the ICS's sign-in page. For example, you might configure the ICS to check whether or not the user's computer is running a particular version of Norton Antivirus. In the event it is not running, you can determine that the user's computer is unsecure and disable access to the ICS's sign-in page until the user has updated the computer's antivirus software.

2. After the user has successfully accessed the ICS's sign-in page, realm-level checks are performed to determine whether the ICS's end-user home page is accessed. The most common realm-level check is performed by an authentication server. The server determines whether the user enters a valid username and password. You can perform other types of realm-level checks such as checking if the user's IP address is in your network or that the user is using the web browser type that you have specified.
3. If the user does not get through the realm-level checks that are specified, the user is not allowed to sign in, or a "stripped down" version of the home page is displayed. Generally, this stripped-down version contains significantly less functionality than what is available to your standard users because the user has not passed all the authentication criteria. The ICS provides extremely flexible policy definitions, enabling you to dynamically alter end-user resource access based on corporate security policies.
4. After the ICS successfully assigns a user to a realm, the appliance maps the user to a role based on your selection criteria. A role specifies which access mechanisms a selected group of users can access. It also controls session and UI options for that group of users. You can use a wide variety of criteria to map users to roles. For example, you can map users to different roles based on endpoint security checks or attributes obtained from an LDAP server or client-side certificate.
5. In most cases, a user's role assignments control which individual resources the user can access. For example, you might configure access to your company's Intranet page using a web resource profile and then specify that all members of the Employees role can access that resource.
6. However, you can choose to further fine tune access to individual resources. For example, you may enable members having the Employees role to access your company's Intranet (as described earlier), also add a resource policy detailed rule that requires users to meet additional criteria to access the resource. An additional example would be, you may require users to be members of the employees' role and to sign into the device during business hours to access your company Intranet.

Integration between Ivanti Connect Secure and the Resources It Intermediates

In a typical configuration, you can add bookmarks directly to the ICS's end-user home page. The bookmarks that you add are links to the resources that you configure the ICS to intermediate. Adding these bookmarks enables the users to sign into the Ivanti Connect Secure and find a consolidated list of resources available for them. Within this typical configuration, you can streamline the integration between the ICS and the intermediated resources by enabling single sign-on (SSO).

SSO is a process that allows pre-authenticated users to access other applications or resources that are protected by another access management system without having to re-enter their credentials. During system configuration, you can enable SSO by specifying user credentials that you want the ICS to pass to the intermediated resources. Alternatively, if you do not want to centralize user resources on the ICS's end-user home page, you can create links to the intermediated resources from another web page.

To cite an example, you can configure bookmarks on Ivanti Connect Secure, and then add links to those bookmarks from your company's Intranet. Your users can then sign into your company's Intranet and click the links there to access the intermediated resources without going through the ICS's home page. As with standard Ivanti Connect Secure bookmarks, you can enable SSO for these external links.

Using Host Checker to Protect from Threats

The Host Checker feature in Ivanti Connect Secure protects the ICS against viruses, attacks, and other security concerns. Host Checker performs security checks on the clients that connect to the device.

Host Checker can:

- Verify if the end-user system contains up-to-date antivirus software, firewalls, critical software hotfixes, and other applications that protect your users' computers.
- Enable or deny users' access to the ICS's sign-in pages, realms, roles, and resources based on the results that Host Checker returns. Alternatively, you can display the recovery instructions to users, so they can bring their computers into compliance.
- Secure your network from hostile outside intrusion by integrating your device with a Juniper Networks Intrusion Detection. You can use IDP devices to detect and block most network worms based on software vulnerabilities, non-file-based Trojan horses, the effects of Spyware, Adware, and Key Loggers, many types of malware, and zero-day attacks with anomaly detection.

Providing Redundancy in the Ivanti Connect Secure Environment

The Clustering feature in Ivanti Connect Secure ensures redundancy in your environment. With this feature you can:

- Deploy two or more appliances as a cluster, ensuring no user downtime in the rare event of failure and stateful peering that synchronizes user settings, ICS settings, and user session data.
- Support active/passive or active/active configurations across a LAN.

- In Active/Passive mode, one device actively serves user requests while the other device runs passively in the background to synchronize state data. If the active device goes offline, the passive device automatically starts servicing user requests.
- In active/active mode, all the machines in the cluster actively handle user requests sent by an external load balancer. The load balancer hosts the cluster VIP and routes user requests to a device defined in its cluster group based on source-IP routing. If a device goes offline, the load balancer adjusts the load on the other active device.

Ivanti supports two-node on ICS active-active cluster and standalone nodes with load balancer. Although four-node active-active cluster is typically known to work in some environments, using mode nodes in the cluster results in the overhead of the cluster operations like session sync and lesser throughput. To get the optimal performance from the gateways Ivanti recommends using two-node clusters or standalone nodes behind a load balancer.



In a well-connected campus network, where the connectivity is more LAN-like than WAN-like, the Ivanti Connect Secure can be clustered in separate buildings.

Customizing the Interface to Match a Company's Look-and-Feel

Ivanti Connect Secure enables you to customize a variety of elements in the end-user interface.

You can use the customization features to:

- Update the look-and-feel of the ICS's end-user console, so it will resemble one of your company's standard web pages or applications.
- Modify the headers, background colors, and logos that display in the sign-in page and end-user console to match your company's style.
- Order the bookmark display at the end user help system.
- Display the end-user home page to users (either in standard or customized form), and then choose to redirect users to a different page (such as your company's Intranet) when users first sign into the ICS console. On choosing to use this option, you may want to add links to your ICS's bookmarks on the new page.
- Configure custom sign-in pages through the ICS's admin console. The custom sign-in pages feature does not limit the number of customizations you can make to your pages. Using this feature, you can use an HTML editor to develop a sign-in page that exactly matches your specifications.

Supporting Users on Different Devices to Access Ivanti Connect Secure

Ivanti Connect Secure is accessed from standard workstations and kiosks running on Windows, Mac OSX, and Linux operating systems. End users can also access the ICS from connected Smart Phones and Tablets.

When a user connects from a Smart Phone or a Tablet, the ICS determines which pages and functionality to display based on settings that you configure.

For more information about specifying which pages get displayed on different devices, see the *Ivanti Connect Secure Supported Platform Guide*.

Providing Secure Access for International Users

Ivanti Connect Secure supports localization to include English (US), French, German, Spanish, Simplified Chinese, Traditional Chinese, Japanese, Korean, Italian, and Polish. When users sign into the device, it automatically detects the correct language to display based on the web browser setting. Alternatively, you can use end-user localization and custom sign-in page options to manually specify the language that you want to display to your end users.

Configuring Ivanti Connect Secure

The following basic steps need to be completed to enable users to start using Ivanti Connect Secure.

1. Plug in the Ivanti Connect Secure device and connect it to your network. Configure the initial system and network settings (see the PSA Series Hardware Guide for more information).
2. When you first sign into the admin console, an initial configuration task guide display, to walk you through the upgrade and installation of product licenses process. To view the configuration task guide, click **Guidance** in the upper right corner of the admin console.
3. Set the system date and time, upgrade to the latest service package, and install your product licenses.
4. Followed by the installation of product licenses, use the following steps to set up your access management framework to enable users to authenticate and access resources.



Create a test scenario to familiarize yourself with the process (see [Creating a Test Scenario to Learn Concepts and Best Practices](#) for more information).

5. Define an authentication server that verifies user names and passwords.

6. Create the user roles that enable access mechanisms, session options, and UI options for user groups.
7. Create a user authentication realm that specifies the conditions that users must meet to sign into the device.
8. Define a sign-in policy that specifies the URL that users must access to sign into the device and the page that they see when they sign in.
9. Create resource profiles that control access to resources, specify which user roles can access them, and include bookmarks that link to the resources.

After completing the basic steps, your system is ready for use. You can start using it as it is or configure additional advanced features such as endpoint defense and clustering.

Introducing the Pulse Clients

The Ivanti Connect Secure gateway is the server component of a larger client-server solution. Ivanti allows many different clients to provide an array of secure-connectivity services to end users. In general, Pulse clients can be divided into three groups:

- Desktop Clients
- Mobile Clients
- Integrated Clients

Each of these clients are described below:

Desktop Clients

Pulse desktop clients are fully-featured secure-connectivity clients that can be deployed either directly from the Ivanti Connect Secure gateway or via other third-party software distribution mechanisms (e.g., SMS). The Pulse desktop clients support Windows and Mac OSX.

The Windows desktop client provides VPN, Host Checker, and Layer-2 (NAC) functionality, whereas the OSX desktop client provides VPN and Host Checker functionality. The Pulse desktop clients can be downloaded from my.pulsesecure.net without having to download the Ivanti Connect Secure gateway packages. Refer to the *Pulse Desktop Clients* documentation for details on desktop clients.

Mobile Clients

Pulse mobile clients differ from the desktop clients in that they are made available through App Stores (rather than hosted on the Ivanti Connect Secure gateway). Ivanti offers mobile clients for iOS, Android, Google Chrome OS, and Windows (the Windows mobile client is also called the "Universal App").

Mobile clients are designed to be lightweight and work tightly within the "sandboxes" provided by the mobile operating systems. The exact functionality of each mobile client varies according to the operating system, so, refer to the *Pulse Mobile Client* documentation for details on the capability of each mobile client.

Integrated Clients

There are many clients that are integrated directly into the Ivanti Connect Secure gateway. They are deployed by the ICS gateway and cannot be acquired independently from the ICS gateway. For the most part, these integrated clients are accessed by end users via a web browser connected to the ICS gateway.

User Verification and Key Concepts

User verification is the process that is supported to identifying a user, authorize and then determine whether a user can take specific actions. The following sections describe the concept and steps behind how user role, sign-in policy and authentication work in Ivanti Connect Secure.

Verifying User Accessibility

Before you access your device, you need to create a user account in the system authentication server for verifying the user accessibility. After creating the account through the admin console, sign in as the user on the user sign-in page.

To verify user accessibility:

1. From the admin console, choose **Authentication > Auth. Servers**.
2. Select the **System Local** link.
3. Select the **Users** tab.
4. Click **New**.
5. Type **testuser1** as the username and enter a password, and then click **Save Changes**. The testuser1 account is now created.
6. Use another browser window to enter the machine's URL to access the user sign-in page. The URL is in the format: `https://a.b.c.d`, where a.b.c.d is the machine IP address you entered in the serial console when you initially configured your device.
7. Click **Yes** when prompted with the security alert to proceed without a signed certificate. The user sign-in page appears, indicating that you have successfully connected to your device.
8. Enter the username and password you created for the user account and then click **Sign In** to access the home page for users.
9. Enter the URL to an internal web server in the **Address** box and click **Browse**. The ICS opens the web page in the same browser window, so to return to the ICS home page, click the center button on the toolbar that appears on the target web page.
10. Enter the URL to your external corporate site on the ICS home page, and click **Browse**. A web page opens in the same browser window, so use the button on the toolbar to return to the ICS home page.

11. Click **Browsing > Windows Files** on the ICS home page to browse through available Windows file shares.

Creating a Test Scenario to Learn Concepts and Best Practices

The ICS provides a flexible access management system that makes it easy to customize a user's remote access experience with roles, resource policies, authentication servers, authentication realms, and sign-in policies.

To enable you to quickly begin working with these entities, your device ships with ICS defaults for each entity that you will work with. You can create each access management entity by performing the tasks in the following sub-sections.

The ICS supports two types of users:

- **Administrators**-An administrator is a person who may view or modify ICS's configuration settings. You create the first administrator account through the serial console.
- **Users**-A user is a person who uses the ICS to gain access to corporate resources as configured by an administrator.

Defining a User Role

Your device is preconfigured with one user role called "Users." This predefined role enables the web and file browsing access features, enabling any user mapped to the Users role to access the Internet, corporate web servers, and any available Windows file servers.

You can view this role on the User Roles page after you enable an access feature for a role, configure the appropriate corresponding options that are accessible from the access feature's configuration tab.

To define a user role:

1. In the admin console, choose **Users > User Roles**.
2. Click **New Role**.
3. Enter Test Role in the **Name** box and then click **Save Changes**.
4. Wait for the ICS to display the Test Role page with the **General** tab and **Overview** link selected.
5. Select the **Web** check box under Access features and then click **Save Changes**.
6. Select **Web > Options**.
7. Select **User can type URLs in the IVE browse bar** check box, and then click **Save Changes**.

After completing these steps, you have defined a user role. When you create resource profiles, you can apply them to this role. You can also map users to this role through role mapping rules defined for an authentication realm.


To quickly create a user role that enables web and file browsing, duplicate the Users role, and then enable additional access features as desired.

Defining a Resource Profile


A resource profile is a set of configuration options that contain all the resource policies, role assignments, and end-user bookmarks required to provide access to an individual resource.

Within a resource profile, a resource policy specifies the resources to which the policy applies (such as URLs, servers, and files) and whether the ICS grants access to a resource or performs an action. Note that the ICS is preconfigured with two types of resource policies:

Web Access - The predefined web Access resource policy, Initial Policy for Local Resources, allows access only to hosts belonging to domains within the secured network.

-
- From release 9.1R3, for a fresh installation, this predefined "Initial Policy for Local Resources" policy is in "Deny" state by default.
-  • From 8.3R1 onwards, to allow access to IPv6 hosts belonging to domains within the secured network, add the [fd00::/8]:*/* resource to the predefined Web Access resource policy, if not present already.
-

Windows Access - The predefined Windows Access resource policy enables all users mapped to the Users role to access all corporate Windows file servers. By default, this resource policy applies to the Users role.

-
- From release 9.1R3, for a fresh installation, this predefined "Initial File Browsing Policy" is in "Deny" state by default.
-  • Delete the Windows Access resource policies if you are concerned about users having access to all your web and file content.
-

To define a resource profile:

1. In the admin console, choose **Users > Resource Profiles > Web**.
2. Click **New Profile**.

The **Web Applications Resource Profile** page appears.

3. Fill in the following information:

1. In the **Type** box, keep the default option (**Custom**).
2. In the **Name** box, type **Test Web Access**.
3. In the **Base URL** box, type <http://www.google.com>
4. Under **Autopolicy: Web Access Control**, select the check box next to the default policy (http://www.google.com:80/*) and choose **Delete**.
5. In the **Resource** box, type <http://www.google.com>, select **Deny** from the **Action** list, and click **Add**.
6. Click **Save** and **Continue**. The Test Web Access page appears.
7. Click the **Roles** tab.
8. Select **Test Role** in the **Available Roles** box and click **Add** to move it to the **Selected Roles** box.
9. Click **Save Changes**.

The ICS adds **Test Web** Access to the web Application Resource Policies page and automatically creates a corresponding bookmark that links to google.com.

After completing these steps, you have configured a web Access Resource profile. Even though the ICS comes with a resource policy that enables access to all web resources, users mapped to Test

Role are still prohibited from accessing <http://www.google.com>. These users are denied access because the auto policy you created during the resource profile configuration takes precedence over the default web access policy that comes with the ICS.

Defining an Authentication Server

An authentication server is a database that stores user credentials - username and password - and typically group and attribute information. When a user signs into the host, the user specifies an authentication realm, which is associated with an authentication server. The ICS forwards the user's credentials to this authentication server to verify the user's identity.

The ICS supports the most common authentication servers, including Active Directory, RADIUS, LDAP, RSA ACE/Server, and SAML Server, and enables you to create one or more local databases of users who are authenticated.

The ICS is pre-configured with one local authentication server for users called "System Local." This predefined local authentication server is a system database that enables you to quickly create user accounts for user authentication. This ability provides flexibility for testing purposes and for providing third-party access by eliminating the need to create user accounts in an external authentication server.

You can view the default local authentication server on the Authentication Servers page.



The ICS also supports authorization servers. An authorization server (or directory server) is a database that stores user attribute and group information. You can configure an authentication realm to use a directory server to retrieve user attribute or group information for use in role mapping rules and resource policies.

To define an authentication server:

1. In the admin console, choose **Authentication > Auth. Servers**.
2. Select **Local Authentication** from the **New** list and then click **New Server**.

The New Local **Authentication** page appears.

3. Enter **Test Server** in the **Name** box and then click **Save Changes**.

Wait for the ICS to notify you that the changes are saved, after which additional configuration tabs appear.

4. Click the **Users** tab and then click **New**.

The New Local User page appears.

5. Enter **testuser2** in the **Username** box, enter a password, and then click **Save Changes** to create the user's account in the Test Server authentication server.

After completing these steps, you have created an authentication server that contains one user account. This user can sign in to an authentication realm that uses the Test Server authentication server.

The admin console provides last access statistics for each user account on the respective authentication server pages, on the Users tab under a set of columns titled **Last Sign-in Statistic**. The statistics reported include the last successful sign-in date and time for each user, the user's IP address, and the agent or browser type and version.

Defining an Authentication Realm

An authentication realm is a grouping of authentication resources, including:

- An authentication server, which verifies a user's identity. The ICS forwards credentials submitted on a sign-in page to an authentication server.
- An authentication policy, which specifies realm security requirements that need to be met before the ICS submits credentials to an authentication server for verification.
- A directory server, which is an LDAP server that provides user and group attribute information to the ICS for use in role mapping rules and resource policies (optional).
- Role mapping rules, which are conditions a user must meet for the ICS to map a user to one or more roles. These conditions are based on information returned by the realm's directory server, the person's username, or certificate attributes.

Your ICS is pre-configured with one user realm called "Users." This predefined realm uses the System Local authentication server, an authentication policy that requires a minimum password length of four characters, no directory server, and contains one role mapping rule that maps all users who sign in to the Users realm to the Users role.

The "testuser1" account you created is part of the Users realm, because this account is created in the System Local authentication server. The "testuser2" account you created is not part of the Users realm, because you create the user account in the new "Test Server" authentication server, which is not used by the Users realm.

You can view the default user authentication realm on the User Authentication Realms page.

To define an authentication realm:

1. In the admin console, choose **Users > User Realms**.

The User Authentication Realms page appears.

2. Click **New**.

The New Authentication Realm page appears.

3. Enter **Test Realm** in the **Name** box.
4. Select **Test Server** from the **Authentication** list.
5. Click **Save Changes**.

Wait for the ICS to notify you that the changes are saved and to display the realm's configuration tabs.

6. Click the **Role Mapping** tab if it is not already selected, and then click **New Rule**.

The Role Mapping Rule page appears.

7. Enter **testuser2** in the **text box**.
8. Under “...then assign these roles”, select **Test Role** from the **Available Roles** list and click **Add to move it to the Selected Roles** box.
9. Click **Save Changes**.

After completing these steps, you have finished creating an authentication realm. This realm uses Test Server to authenticate users and a role mapping rule to map testuser2 to Test Role. Because the Test Web Access resource policy applies to Test Role, any user mapped to this role cannot access <http://www.google.com>

Defining a Sign-In Policy

A sign-in policy is a system rule that specifies:

- A URL where a user may sign in to the host.
- A sign-in page to display to the user.
- Whether or not the user needs to type or select an authentication realm to which the ICS submits credentials.
- The authentication realms where the sign-in policy applies.

You can view the default user sign-in policy on the Signing In page.

To define a sign-in policy:

1. In the admin console, choose **Authentication > Signing in > Sign-in Policies**.
The Signing In page appears.
2. Click ***/** under User URLs.
The ***/** page appears.
3. Enter **test** after ***/** in the **Sign-in URL** box.
4. Under Authentication realm, select the **User picks from a list of authentication realms** option button
5. Select **Test Realm** from the **Available Realms** list. Click **Add** to move it to the **Selected Realms** box. (Repeat this process for the Users role if it is not already in the **Selected Realms** box.)
6. Click **Save Changes**.

After completing these steps, you have finished modifying the default users sign-in policy.

Optional Steps

You can perform these following optional steps to define a new sign-in page that is associated with the */test/ sign-in policy.

1. Select **Authentication > Signing In > Sign In Pages**, and then click **New Page**.
2. Enter **Test Sign-in Page** in the **Name** field, type **#FF0000 (red)** in the **Background color** box, and then click **Save Changes**.
3. Select **Authentication > Signing In > Signing In Policies**, and then click **New URL**.

The New Sign-In Policy page appears.

4. Type ***/test/** in the **Sign-in URL** box.
5. Select **Default Sign-in Page** from the **Sign-in Page** list, and click **Save Changes**.
6. Select **Authentication > Signing In > Sign In Policies**, and then click ***/test/ under User URLs**.

The */test/ page appears.

7. Select **Test Sign-in Page from the Sign-in page** list and then click **Save Changes**.

All ICS devices are pre-configured with one sign-in policy that applies to users:

/. This default user sign-in policy (/) specifies that when a user enters the URL to the host, it displays the default sign-in page for the user and requires the user to select an authentication realm (if more than one realm exists). The */ sign-in policy is configured to apply to the Users authentication realm, therefore this sign-in policy does not apply to the authentication realm you created.

Using the Test Scenario

The test scenario enables you to do the following tasks:

- Access the user console using the modified default sign-in policy.
- Sign in as the user created in the Test Server to map to the Test Realm.
- Test your web browsing capabilities, which are dependent upon the proper configuration of Test Role and Test Web Access.

To use the test scenario:

1. In a browser, enter the User URL followed by /test to access the user sign-in page. The URL is in the format: `https://a.b.c.d/test`, where a.b.c.d is the machine IP address you entered in the serial console during initial configuration.
2. Click **Yes** when prompted with the security alert to proceed without a signed certificate. If the user sign-in page appears, you have successfully connected to your device.



If you performed the optional configuration steps in "Defining a Sign-In Policy", the header color is red.

3. Enter the username and password you created for the user account in Test Server, type **Test Realm** in the **Realm** box, and then click **Sign In** to access the ICS home page for users.

The ICS forwards the credentials to Test Realm, which is configured to use Test Server. Upon successful verification by this authentication server, the ICS processes the role mapping rule defined for Test Realm, which maps testuser2 to Test Role. Test Role enables web browsing for users.

4. In the browser Address bar, enter the URL to your corporate web site and click **Browse**. The web page opens in the same browser window, so to return to the ICS home page, click the Home icon in the browsing toolbar that appears on the target Web page.
5. On the ICS home page, type `www.google.com` and click **Browse**. An error message appears because the Test Web Access resource policy denies access to this site for users mapped to Test Role.
6. Return to the ICS home page, click **Sign Out**, and then return to the user sign-in page.
7. Enter the credentials for testuser1, specify the Users realm, and then click **Sign In**.
8. On the ICS home page, type `www.google.com` and click **Browse**. The web page opens in the same browser window.
 - The test scenario demonstrates the basic access management mechanisms. You can create very sophisticated role mapping rules and resource policies that control user access depending on factors such as a realm's authentication policy, a user's group membership, and other variables.
 - To learn more about access management, we recommend that you take a few minutes to review the Online Help to familiarize yourself with its contents.

- When you configure your device for your enterprise, we recommend that you perform user access configuration. Before you make your device available from external locations, we recommend that you import a signed digital certificate from a trusted certificate authority (CA).

Default Settings for Administrators

Just like for users, the ICS provides default settings that enable you to quickly configure accounts for administrators. This list summarizes the ICS default settings for administrators:

- **Administrator roles** - There are two built-in administrator roles.
 - **Administrators** - This built-in role permits administrators to manage all aspects of the device. The administrator user you create through the serial console is mapped to this role.
 - **Read-Only Administrators** - This built-in role permits users mapped to the role to view (but not configure) all settings. You need to map administrators to this role if you want to restrict their access.
- Administrator local authentication server is a database that stores administrator accounts. You create the first administrator account in this server through the serial console. (All administrator accounts created through the serial console are added to this server.) You cannot delete this local server.
- Admin Users authentication realm uses the default Administrators local authentication server, an authentication policy that requires a minimum password length of 10 characters, no directory server, and one role mapping rule that maps all users who sign in to the Admin Users realm to the Administrators role. The administrator account you create through the serial console is part of the Admin Users realm.



From 9.1R3 release onwards, minimum password length should be 10 characters when deploying ICS/IPS VA on Azure Cloud, AWS Cloud, or AliCloud.

- `*/admin` sign-in policy is the default administrator sign-in policy. The `*/admin` specifies that when a user enters the URL to the host followed by `/admin`, the ICS displays the default sign-in page for administrators. This policy also requires the administrator to select an authentication realm (if more than one realm exists).
The `*/admin` sign-in policy is configured to apply to the Admin Users authentication realm, therefore this sign-in policy applies to the administrator account you create through the serial console.

General Access Management

Access Management Overview

The system enables you to secure your company resources using authentication realms, user roles, and resource policies. These three levels of accessibility allow you to control access from a very broad level (controlling who may sign into the device) down to a very granular level (controlling which authenticated users may access a particular URL or file). You can specify security requirements that users must meet to sign in to the device, to gain access to features, and even to access specific URLs, files, and other server resources. The system enforces the policies, rules and restrictions, and conditions that you configure to prevent users from connecting to or downloading unauthorized resources and content.

To permit endpoints that are not directly connected to a Ivanti security device to access resources behind the firewall, you can configure a Policy Secure device to provision shared user sessions from any number of different Ivanti Connect Secure devices and Infranet Controllers. IF-MAP Federation allows users to access resources protected by any number of Ivanti firewalls (Infranet Enforcers) without requiring additional authentication.

The access management framework is available on all Ivanti Connect Secure products. The access management features, including realms, roles, resource policies, and servers, are the base of the platform on which all Ivanti Connect Secure products are built.

Policies, Rules & Restrictions, and Conditions Overview

The system enables you to secure your company resources using authentication realms, user roles, and resource policies. These three levels of accessibility allow you to control access from a very broad level (controlling who may sign into the device) down to a very granular level (controlling which authenticated users may access a particular URL or file).

Accessing Authentication Realms

Resource accessibility begins with the authentication realm. An authentication realm is a grouping of authentication resources, including:

- **An authentication server** - verifies that the user is who one claims to be. The system forwards credentials that a user submits on a sign-in page to an authentication server.
- **An authentication policy** - specifies realm security requirements that need to be met before the system submits a user's credentials to an authentication server for verification.

- **A directory server** - specifies an LDAP server that provides user and group information to the system that it uses to map users to one or more user roles.
- **Role mapping rules** - specifies the conditions a user must meet for the system to map the user to one or more user roles. These conditions are based on either user information returned by the realm's directory server or the user's username.
 - You can associate one or more authentication realms with the sign-in page. When more than one realm exists for a sign-in page, a user must specify a realm before submitting one's credentials. When a user submits their credentials, the system checks the authentication policy defined for the chosen realm. The user must meet the security requirements you define for a realm's authentication policy or else the system does not forward the user's credentials to the authentication server.
 - At the realm level, you can specify security requirements based on various elements such as the user's source IP address or the possession of a client-side certificate. If the user meets the requirements specified by the realm's authentication policy, the system forwards the user's credentials to the appropriate authentication server. If this server successfully authenticates the user, then the system evaluates the role mapping rules defined for the realm to determine which roles to assign to the user.

Accessing User Roles

A role is a defined entity that specifies session properties for users who are mapped to the role. These session properties include information such as session time-outs and enabled access features. A role's configuration serves as the second level of resource access control. Not only does a role specify the access mechanisms available to a user, but you can also specify restrictions with which users must comply before they are mapped to a role.

At the role level, you can specify security requirements based on elements such as the user's source IP address and possession of a client-side certificate. If the user meets the requirements specified either by a role mapping rule or a role's restrictions, then the system maps the user to the role. When a user makes a request to the backend resources available to the role, the system evaluates the corresponding access feature resource policies.

Note that you may specify security requirements for a role in two places in the role mapping rules of an authentication realm (using custom expressions) or by defining restrictions in the role definition. The system evaluates the requirements specified in both areas to make sure the user complies before it maps the user to a role.

Accessing Resource Policies

A resource policy is a set of resource names (such as URLs, hostnames, and IP address/netmask combinations) to which you grant or deny access or other resource-specific actions, such as rewriting and caching. A resource policy serves as the third level of resource access control. While a role may grant access to certain types of access features and resources (such as bookmarks and applications), whether or not a user can access a specific resource is controlled by resource policies. These policies may even specify conditions that, if met, either deny or grant user access to a server share or file. These conditions may be based on security requirements that you specify. The user must meet these security requirements or else the system does not process the user's request.

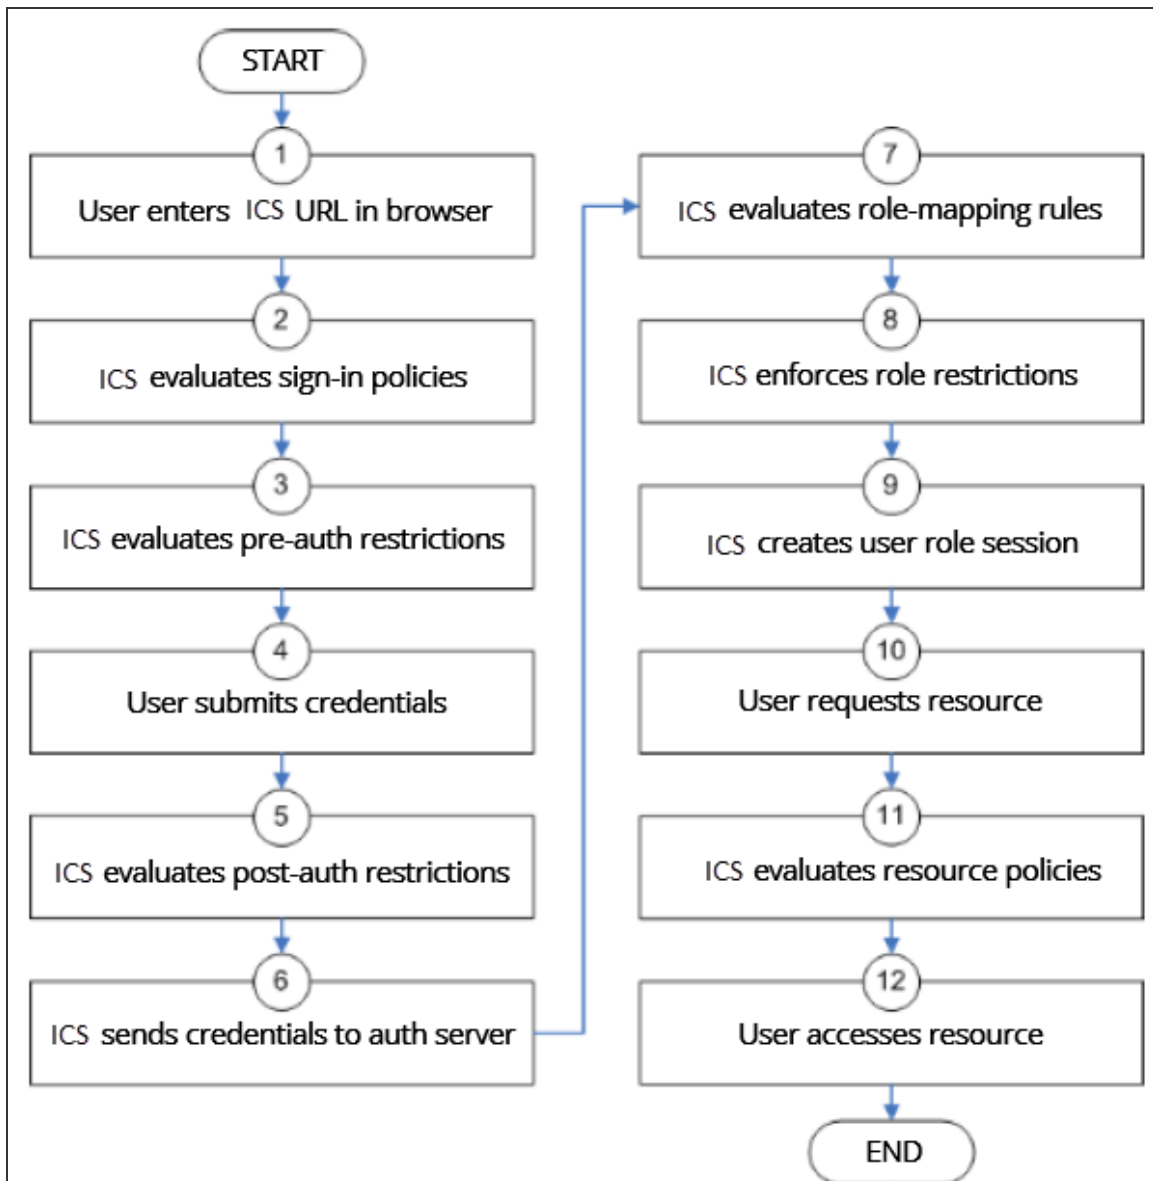
At the resource level, you can specify security requirements based elements such as the user's source IP address or possession of a client-side certificate. If the user meets the requirements specified by a resource policy's conditions, then the system either denies or grants access to the requested resource. You may enable Web access at the role level, for example, and a user mapped to the role may make a Web request. You may also configure a Web resource policy to deny requests to a particular URL or path when Host Checker finds an unacceptable file on the user's machine. In this scenario, the system checks to see if Host Checker is running and indicates that the user's machine complies with the required Host Checker policy. If the user's machine complies, meaning the unacceptable file is not found, then the system grants the user access to the requested Web resource.

Note that you can create separate resource policies, or you can create automatic resource policies (called autopolicies) during resource profile configuration (recommended).

Policies, Rules & Restrictions, and Conditions Evaluation

The following figure illustrates the access management security checks that the system performs when a user tries to access resources through the device. A detailed description of each step follows the diagram.

The following figure depicts the Security Checks Performed During a User Session:



1. The user enters the URL of the device end-user console (such as <http://employees.yourcompany.com/marketing>) in a web browser.
2. The system evaluates its sign-in policies (starting with the administrator URLs and continuing to user URLs) until it matches the hostname entered by the user.

3. The system evaluates pre-authentication restrictions and determines if the user's system passes host checks and other requirements. If the pre-authentication checks fail, the system denies the user access. If the checks pass, the system prompts the user to enter the username and password for the realms whose preauthentication checks succeeded. (If required by the realm, the system prompts the user to enter two sets of credentials.) If more than one realm exists, the user must enter a realm or choose one from a list.
4. The system evaluates the post-authentication restrictions and determines whether the user's password conforms to specified limits and requirements. If the postauthentication checks fail, the system denies the user access. If the checks pass, the system passes the user's credentials to the realm's authentication server.
5. The system forwards the user's username and password to the authentication server, which returns success or failure. (A RADIUS returns attributes for the system to use in role mapping.) If the authentication server returns failure, the system denies the user access. If the server returns success, the system stores the user's credentials. If the realm has a separate LDAP authorization server, the system also queries the LDAP server for attribute and group information and saves the information returned by LDAP. If the realm includes a secondary authentication server, the system repeats this process with the secondary server.
6. The system evaluates the realm's role mapping rules and determines the roles for which the user is eligible. The system determines eligibility using information from the LDAP or RADIUS server or the user's username.
7. The system evaluates the restrictions of the eligible roles, enabling the user to access those roles whose restrictions the user's computer meets. Restrictions may include source IP, browser type, client-side certificate, and Host Checker.
8. The system creates a "session role," determining the user's session permissions. If you enable permissive merging, the system determines session permissions by merging all valid roles and granting the allowed resources from each valid role. If you disable merging, the system assigns the user to the first role to which he is mapped.
9. When the user requests a resource, the system checks whether the corresponding access feature is enabled for the session user role. If not, the system denies the user access. If the access feature is enabled, the evaluates resource policies.
10. The system evaluates resource profiles and policies related to the user's request, sequentially processing each until it finds the profile or policy whose resource list and designated roles match the user's request. The system denies user access to the resource if specified by the profile or policy. Otherwise, the system intermediates the user request if the profile or policy enables access.

11. The system intermediates the user request, forwarding the user's request and credentials (if necessary) to the appropriate server. Then, the system forwards the server's response to the user.
12. The user accesses the requested resource or application server. The user session ends when the user signs out or the session times out due to time limits or inactivity. The system may also force the user out if the session if you enable dynamic policy evaluation and the user fails a policy.
13. The user can perform realm, role mappings and create rules based on the Enhanced Key Usage (EKU) attribute in the certificates. This attribute can be parsed in certificates to create realm restrictions, role restrictions and role mapping based on rules that contained this attribute. Also, this is supported for custom expressions. The Enhanced Key Usage has 2 parts - The EKU Text and the EKU OID. The EKU text has information about the enhanced key usage - for example - "smart card logon", "wireless", "TLS Web Server Authentication", "E-mail Protection", "TLS Web Client Authentication" and so on. The OID is an identifier for this attribute and is a dotted number representation. The restrictions and role mappings can be done on either the text or the OID.



If you enable dynamic policy evaluation, the system performs additional checks beyond the ones mentioned here.

Dynamic Policy Evaluation

Dynamic policy evaluation allows you to automatically or manually refresh the assigned roles of users by evaluating a realm's authentication policy, role mappings, role restrictions, and resource policies. When the system performs a dynamic evaluation, it checks whether the client's status has changed. (For instance, the client's Host Checker status may have changed. Or, if the user is roaming, the computer's IP address may have changed.) If the status has changed, the system enables or denies the user access to the dependent realms, roles, or resource policies accordingly.

The system does not check for changes in user attributes from a RADIUS, or LDAP server when performing dynamic policy evaluation. Instead, the system re-evaluates rules and policies based on the original user attributes that it obtained when the user signed into the device.

Understanding Dynamic Policy Evaluation

Please note the following about Dynamic Policy Evaluation:

- Clients that use Network Communications Protocol (NCP) do not honor policy changes. This includes PSAM and WTS/CTS.

- PSAM establishes a new NCP tunnel when the protected application opens a new connection, so PSAM establishes new NCP connections frequently. This means PSAM gets the new policy frequently.
- WTS has a persistent NCP tunnel so it does not get policy changes until the user disconnects and then reconnects.



Because the system evaluates Web and Files resource policies whenever the user makes a request for a resource, dynamic policy evaluation is unnecessary for Web and Files.

If the system determines after a dynamic policy evaluation that a user no longer meets the security requirements of a policy or role, the system terminates the connection immediately with the user. The user may see the closing of a TCP or application connection, or the termination of a user session for VPN Tunneling, Secure Application Manager, Terminal. The user must take the necessary steps to meet the security requirements of the policy or role, and then sign into the system again.

The system logs information about policy evaluation and changes in roles or access in the Event log.

Understanding Standard Policy Evaluation

If you do not use dynamic policy evaluation, the system evaluates policies and roles only when the following events occur:

- When the user first tries to access the system sign-in page, the system evaluates the Host Checker policies (if any) for a realm.
- Immediately after the user's initial authentication, the system evaluates the user's realm restrictions in the authentication policy, role mapping rules, and role restrictions.
- When the user makes a request for a resource, the system evaluates resource access policies to determine if the associated role is allowed to access the resource.
- When the Host Checker status of the user's machine changes, the system evaluates the Host Checker policies (if any) for the role.

If you do not use dynamic policy evaluation and you make changes to an authentication policy, role mapping rules, role restrictions, or resource policies, the system enforces those changes only when the events described above occur.

If you use dynamic policy evaluation, the system enforces changes when the events described above occur, and it also enforces changes at the times you specify.

Enabling Dynamic Policy Evaluation

You can use dynamic policy evaluation in the following ways:

- **Evaluate all signed-in users in a realm** - You can automatically or manually refresh the roles of all currently signed-in users of a realm by using the General tab of the Administrators > Admin Realms > Select Realm or Users > User Realms > Select Realm page. You can trigger the system to perform a dynamic policy evaluation at the realm level based on:
 - **An automatic timer** - You can specify a refresh interval that determines how often the system performs an automatic policy evaluation of all currently signed-in realm users, such as every 30 minutes. When using the refresh interval, you can also fine tune the system performance by specifying whether or not you want to refresh roles and resource policies as well as the authentication policy, role mapping rules, and role restrictions.
 - **On-demand** - At any time, you can manually evaluate the authentication policy, role mapping rules, role restrictions, and resource policies of all currently signed-in realm users. This technique is especially useful if you make changes to an authentication policy, role mapping rules, role restrictions, or resource policies and you want to immediately refresh the roles of a realm's users.
- **Evaluate all signed-in users in all realms** - At any time, you can manually refresh the roles of all currently signed-in users in all realms by using settings in the System > Status > Active Users page.
- **Evaluate individual users** - You can automatically refresh the roles of individual users by enabling dynamic policy evaluation for Host Checker on the Authentication > Endpoint Security > Host Checker page. Host Checker can trigger the system to evaluate resource policies whenever a user's Host Checker status changes. (If you do not enable dynamic policy evaluation for Host Checker, the system does not evaluate resource policies, but it does evaluate the authentication policy, role mapping rules, and role restrictions whenever a user's Host Checker status changes.)

Specifying Source IP Access Restrictions

This topic describes options to enforce source IP restrictions for access to the corporate network or intranet resources.

About Source IP Restrictions

You can enforce access rules based on the source IP address of the request. You can configure rules related to sign-in, role-mapping, and resource access.

At the realm level, you can add source IP rules to the realms associated with sign-in pages. The user must sign in from a host with an IP address that is allowed by the source IP requirements for the authentication realm. If the source IP policy does not allow the host to access the realm, the system does not forward the user's credentials to the authentication server, and the user is denied access. You can set up multiple rules. For example, you can deny access to all users on a wireless network (10.64.4.100), and allow access to all other network users (0.0.0.0).

At the user role level, you can add source IP rules to the criteria that determine user role membership. If the source IP rule disqualifies a user from a role, subsequent role mapping rules are consulted.

In resource policies, you can add allow/deny rules based on source IP.

Specifying Source IP Restrictions at the Realm Level

To specify source IP restrictions:

1. Navigate to the administrator or user realm you want to configure:
 - **Administrators > Admin Realms > Realm**
 - **Users > User Realms > Realm**
2. Select **Authentication Policy > Source IP** to display the Source IP policy configuration page.
3. Choose one of the following options:
 - **Allow users to sign in from any IP address** - Essentially, this option turns off source IP restrictions.
 - **Allow or deny users from the following IP addresses** - Specifies source IP restrictions. If you select this option, use the policy table controls to create source IP rules.
 - **Allow or deny users from the following locations** - Specifies source IP restrictions based on the geographical location of the user. Use the table controls to create location restriction rules. Ensure UEBA package is uploaded on the ICS for this feature to work.



IP address and Location restrictions cannot be applied simultaneously.

4. Add a rule to the table:

1. Use the text boxes to specify source IP address match criteria:
 - For IPv4 clients, enter IPv4 address and netmask pairs.
 - For IPv6 clients, enter IPv6 address and prefix length pairs.
2. Use the **Allow** and **Deny** option buttons to specify the action when a rule matches.
3. Click **Add** to add the rule to the table.
4. Use the selection box and arrow buttons to order the list. Move the highest priority restrictions to the top of the list. For example, to deny access to all users on a wireless network (10.64.4.100) and allow access to all other network users (0.0.0.0), move the wireless network address (10.64.4.100) to the top of the list and move the (0.0.0.0) network below the wireless network.
5. For administrator realms, select the ports where the administrator can log in (internal, external, and management). On virtual appliances that use traffic segregation, administrators can log in on the management port on the Default Network or Administrative Network (see [Using the Traffic Segregation Feature](#)). If necessary, click **External Port or Management Port** to enable the port.
6. Save the configuration.

Specifying Source IP Restrictions at the Role Level

To specify source IP restrictions:

1. Navigate to the administrator or user role you want to configure:
 - **Administrators > Admin Roles > Role**
 - **Users > User Roles > Role**
2. Select **General > Restrictions > Source IP** to display the Source IP policy configuration page.
3. Choose one of the following options:
 - **Allow users to sign in from any IP address** - Essentially, this option turns off source IP restrictions.
 - **Allow or deny users from the following IP addresses** - Specifies source IP restrictions. If you select this option, use the policy table controls to create source IP rules.
4. Add a rule to the table:

5. Use the text boxes to specify source IP address match criteria:
 - For IPv4 clients, enter IPv4 address and netmask pairs.
 - For IPv6 clients, enter IPv6 address and prefix length pairs.
6. Use the **Allow** and **Deny** option buttons to specify the action when a rule matches.
7. Click **Add** to add the rule to the table.
8. Use the selection box and arrow buttons to order the list. Move the highest priority restrictions to the top of the list. For example, to deny access to all users on a wireless network (10.64.4.100) and allow access to all other network users (0.0.0.0), move the wireless network address (10.64.4.100) to the top of the list and move the (0.0.0.0) network below the wireless network.
9. Save the configuration.

Specifying Source IP Restrictions in Resource Policies

A third way to use source IP restrictions is by creating custom rules in resource policies. The action for custom rules is either allow or deny. The match criteria include resources and conditions. One of the conditions you can set is source IP, so you can enforce source IP restrictions through resource policies. For example:

1. Navigate to **Users > Resource Policies**.
2. Select a policy. Click **Web Access Policies**, for example, to display its policies list.
3. Click the **Initial Policy for Local Resources** policy to edit it.
4. Click the **Detailed Rules** tab.
5. Under Conditions, expand the **Prebuilt Conditions** list, expand the **SourceIPStr** selections, select one of the example expressions, such as **SourceIPStr = "192.168.10.0/24"** or **SourceIPStr = "2001:DB8::15"**, and click **Insert Expression** to add the string to the **Conditions** box.
6. Modify the IP address. In other words, replace **192.168.10.0/24** with an IPv4 address / netmask pair; replace **2001:DB8::15** with an IPv6 address.
7. Specify the other match condition (resource) and specify the action (allow or deny).
8. Save the configuration.

Specifying Browser Access Restrictions

Use a browser restriction to control from which Web browsers users can access a system sign-in page or be mapped to a role. If a user tries to sign in to the device using an unsupported browser, the sign-in attempt fails. This feature also enables you to ensure that users sign in to the device from browsers that are compatible with corporate applications or are approved by corporate security policies.

You can restrict system and resource access by browser-type:

- **When administrators or users try to sign in to Ivanti Connect Secure** - The user must sign in from a browser whose user-agent string meets the specified user-agent string pattern requirements for the selected authentication realm. If the realm "allows" the browser's user-agent string, then the system submits the user's credentials to the authentication server. If the realm "denies" the browser's user-agent string, then the system does not submit the user's credentials to the authentication server.
- **When administrators or users are mapped to a role** - The authenticated user must be signed in from a browser whose user-agent string meets the specified user-agent string pattern requirements for each role to which the system may map the user. If the user-agent string does not meet the "allowed" or "denied" requirements for a role, then the system does not map the user to that role.
- **When users request a resource** - The authenticated, authorized user must make a resource request from a browser whose user-agent string meets the specified "allowed" or "denied" requirements for the resource policy corresponding to the user's request. If the user-agent string does not meet the "allowed" or "denied" requirements for a resource, then the system does not allow the user to access the resource.

The browser restrictions feature is not intended as a strict access control since browser user-agent strings can be changed by a technical user. It serves as an advisory access control for normal usage scenarios.

To specify browser restrictions:

1. Select the level at which you want to implement browser restrictions:
 - **Realm level** - Navigate to:
 - **Administrators > Admin Realms > *Select Realm* > Authentication Policy > Browser**
 - **Users > User Realms > *Select Realm* > Authentication Policy > Browser**
 - **Role level** - Navigate to:

- **Administrators > Admin Realms > *Select Realm* > Role Mapping > Select|Create Rule based on Custom Expressions**
- **Administrators > Admin Roles > *Select Role* > General > Restrictions > Browser**
- **Users > User Realms > *Select Realm* > Role Mapping > Select|Create Rule based on Custom Expression**
- **Users > User Roles > *Select Role* > General > Restrictions > Browser**

2. Choose one of the following options:

- **Allow all users matching any user-agent string sent by the browser** - Allows users to access the system or resources using any of the supported Web browsers.
- **Only allow users matching the following User - agent policy**-Allows you to define browser access control rules. To create a rule:

1. For the User-agent string pattern, enter a string in the format

`*browser_string*`

where start (*) is an optional character used to match any character and browser_string is a case-sensitive pattern that must match a substring in the user-agent header sent by the browser. Note that you cannot include escape characters (\) in browser restrictions.

For example, the following is a browser sent user-agent header:

```
Mozilla/5.0 (Windows NT 6.1; WOW64)
```

```
AppleWebKit/537.22 (KHTML, like Gecko)
```

where:

- *Mozilla/5.0* indicates compatibility with the Mozilla rendering engine.
- *(Windows NT 6.1; WOW64)* are the details of the system in which the browser is running.
- *AppleWebKit/537.22* is the platform the browser users.
- *(KHTML, like Gecko)* is the browser platform details.

Using the above example, enter `*Windows NT*` as a string pattern for specifying the Windows NT system. For more details on user-agent strings, see your specific browser's documentation.

2. Select either:

- **Allow** to allow users to use a browser that has a user-agent header containing the <browser_string> substring.
- **Deny** to prevent users from using a browser that has a user-agent header containing the <browser_string> substring.

3. Click **Add**.

3. Click **Save Changes** to save your settings.

Rules are applied in order, so the first matched rule applies.

Literal characters in rules are case sensitive, and spaces are allowed as literal characters.

For example, the string *Netscape* matches any user-agent string that contains the substring Netscape.

The following rule set grants resource access only when users are signed in using Internet Explorer 5.5x or Internet Explorer 6.x. This example takes into account some major non-IE browsers that send the 'MSIE' substring in their user-agent headers:

***Opera*Deny**

***AOL*Deny**

***MSIE 5.5*Allow**

***MSIE 6.*Allow**

Deny

Specifying Certificate Access Restrictions

When you install a client-side certificate on the device through the System > Configuration > Certificates > Trusted Client CAs page of the admin console, you can restrict system and resource access by requiring client-side certificates:

- **When administrators or users try to sign in to Ivanti Connect Secure** - The user must sign in from a machine that possesses the specified client-side certificate (from the proper certificate authority (CA) and possessing any optionally specified field/value pair requirements). If the user's machine does not possess the certificate information required by the realm, the user can access the sign-in page, but once the system determines that the user's browser does not possess the certificate, the system does not submit the user's credentials to the authentication server and the user cannot access features on the device.

To implement certificate restrictions at the realm level, navigate to:

- **Administrators > Admin Realms > *SelectRealm* > Authentication Policy > Certificate**
- **Users > User Realms > *SelectRealm* > Authentication Policy > Certificate**
- **When administrators or users are mapped to a role** - The authenticated user must be signed in from a machine that meets the specified client-side certificate requirements (proper certificate authority (CA) and optionally specified field/value pair requirements) for each role to which the system may map the user. If the user's machine does not possess the certificate information required by a role, then the system does not map the user to that role.
 - **Administrators > Admin Roles > *SelectRole* > General > Restrictions > Certificate**
 - **Users > User Realms > *Select Realm* Role Mapping > *Select|CreateRule* > CustomExpression**
 - **Users > User Roles > *SelectRole* > General > Restrictions > Certificate**
- **When users request a resource** - The authenticated, authorized user must make a resource request from a machine that meets the specified client-side certificate requirements (proper certificate authority (CA) and optionally specified field/value pair requirements) for the resource policy corresponding to the user's request. If the user's machine does not possess the certificate information required by a resource, then the system does not allow the user to access the resource.
 - **Users > Resource Policies > *SelectResource* > *SelectPolicy* > Detailed Rules *Select|CreateRule* > *ConditionField***

The user can perform realm, role mappings and create rules based on the Enhanced Key Usage (EKU) attribute in the certificates. This attribute can be parsed in certificates to create realm restrictions, role restrictions and role mapping based on rules that contained this attribute.



Also, this is supported for custom expressions. The Enhanced Key Usage has two parts - the EKU Text and the EKU OID. The EKU text has information about the enhanced key usage - for example - "smart card logon", "wireless", "TLS Web Server Authentication", "E-mail Protection", "TLS Web Client Authentication" and so on. The OID is an identifier for this attribute and is a dotted number representation. The restrictions and role mappings can be done on either the text or the OID.

Specifying Password Access Restrictions

You can restrict system and resource access by password-length when administrators or users try to sign in to the device. The user must enter a password whose length meets the minimum password-length requirement specified for the realm. Note that local user and administrator records are stored in the system authentication server. This server requires that passwords are a minimum length of 6 characters by default, regardless of the value you specify for the realm's authentication policy.

To specify password restrictions:

1. Select an administrator or user realm for which you want to implement password restrictions.

Navigate to:

- **Administrators > Admin Realms > *Select Realm* > Authentication Policy > Password**
- **Users > User Realms > *Select Realm* > Authentication Policy > Password**

2. Choose one of the following options:

- **Allow all users (passwords of any length)** - Does not apply password length restrictions to users signing in to the device.
- **Only allow users that have passwords of a minimum length** - Requires the user to enter a password with a minimum length of the number specified.



This option is not applicable for IKEv2 users and therefore is not enforced for IKEv2 users.

3. Select **Enable Password Management** if you want to enable password management. You must also configure password management on the authentication server configuration page (local authentication server) or through an LDAP server.

4. If you have enabled a secondary authentication server, specify password length restrictions using the restrictions above as a guideline.
5. Click **Save Changes** to save your settings.

By default, the system requires that user passwords entered on the sign-in page be a minimum of four characters. The authentication server used to validate a user's credentials may require a different minimum length. The local authentication database, for example, requires user passwords to be a minimum length of six characters.

Specifying Session Limits

In addition to the access management options you may specify a limit for concurrent users. A user who enters a URL to one of this realm's sign-in pages must meet any access management and concurrent user requirements specified for the authentication policy before the system presents the sign-in page to the user.

Setting the minimum or maximum setting limit amount allows you to configure realms that are more likely to be available when the device is nearing the amount of licensed users.

Valid numbers for the minimum amount of sessions are between 0 and the license limit. A default of 0 means there is no limits. All of the realms minimum limits can add up to the license limit but cannot exceed it. You cannot modify an existing realm's minimum limit or add a new realm's minimum limit that exceeds the license limit. The maximum limit can be equal to or greater than the minimum limit for a particular realm. Value 0 for maximum limit means no user can log in to the realm.

You can also limit the number of concurrent users per session; a user can have multiple sessions. For example, if a user logs in from two machines in the same realm, an additional session is created. Each session counts towards the user license.

Users who enter through a realm with this feature enabled must have no more than the specified number of sessions open. If the user attempts to open a new session that exceeds the limit, a message appears that denies access or allows the user to continue or cancel.

When considering concurrent users per session, make note of the following:

- All session-related SSO attributes are saved in their respective session in the cache. These attributes are not shared with other sessions.
- All form-related SSO attributes are saved in their respective session in the cache. These attributes are not shared with other sessions.

- All Form-SSO related attributes are saved in their respective session in the cache. The Form SSO state will not be shared with other sessions. The admin configured Form SSO values will be shared across all sessions.
- End-user's home page changes are reflected across all sessions. Any changes to the following will appear in the other concurrent sessions:
 - Bookmarks
 - Panel sorting (Preferences > User Home)
 - E-mail information, Daylight Saving Time (Preferences > General)
 - Autostart Client Application Session (Preferences > Applications)
 - Cached E-mail Info settings (Preferences > Advanced)
 - Delete Cookies (Preferences > Advanced) now has options to let you remove cookies from the current session only or to remove cookies from all sessions.
 - Delete Password (Preferences > Advanced) now has options to let you remove passwords from the current session only or to remove passwords stored by all sessions.
- Host Checker information is saved in each session. They are not shared across concurrent sessions.
- Log messages will contain session identifiers (concatenated at the end of the log message) to differentiate which session the message refers to.
- Each user session maintains its own VPN Tunneling information. This information is not shared between concurrent sessions. However, administrator network connect sessions are shared between concurrent sessions.
- If you log in to the device as administrator, the first session is edit mode. If you log in as an administrator in a concurrent session, that administrator is logged in as read-only mode.
- VPN Tunneling bandwidth allocation is enforced on a per-session basis. For example, if a user is allocated a 1M bandwidth then each user session has a 1M bandwidth. The total bandwidth for this user is the number of sessions of this user times 1M.
- Users can launch terminal services, JSAM or PSAM from any session. Session information is saved per each session; they are not shared across concurrent sessions. Multiple instances of terminal services, JSAM and PSAM cannot be started on the same client.



If you enable the multiple sessions per user feature, IKEv2 clients and VPN Tunneling clients may not be assigned the same IP address. For example, an IKEv2/VPN Tunneling client may be assigned a different VPN Tunneling VIP address each time they connect to the device when the system is obtaining the DHCP addresses from a DHCP server.

Use limits restrictions to set minimum and maximum concurrent users on the realm.

To specify the number of concurrent users limit restrictions:

1. Select an administrator or user realm for which you want to implement limits restrictions.
 - **Administrators > Admin Realms > *SelectRealm* > Authentication Policy > Limits**
 - **Users > User Realms > *SelectRealm* > Authentication Policy > Limits**
2. To limit the number of concurrent users on the realm, select **Limit the number of concurrent users** and then specify limit values for these options:
 - **Guaranteed minimum** - You can specify any number of users between zero (0) and the maximum number of concurrent users defined for the realm, or you can set the number up to the maximum allowed by your license if there is no realm maximum.
 - **Maximum (optional)** - You can specify any number of concurrent users from the minimum number you specified up to the maximum number of licensed users. If you enter a zero (0) into the Maximum field, no users are allowed to log into the realm.
3. Click **Save Changes**.

To specify the number of concurrent users per session limit restriction:

1. Select **Authentication > Signing In > Sign-in Policies**.
2. Select the **Restrict access to administrators only** to immediately terminate all user sessions from all nodes across the cluster. Once enabled, only administrator URLs are accessible across the cluster. Note that Administrators can attempt to sign in even if all rules on this page are disabled.
3. Select the **Enable multiple user sessions** check box to allow users to have multiple concurrent sessions, and specify whether the user can log in when the maximum number of sessions is reached:
 - **Deny any more session from the user**-Displays a message saying the login is denied because it would exceed the maximum number of concurrent sessions.

- **Allow the user to login**-Allows the user to log in. If the Display open user session[s] warning notification option is enabled, the user can select which session to close; otherwise the session that has been idle the longest is closed automatically.
4. Select the **Display open user session[s] warning notification** check box to allow users who have met the maximum session count to close one of their existing sessions before continuing with the current log in. If this option is disabled, the system terminates the session that has been idle the longest. This option applies only if **Enable multiple user sessions** is enabled along with Allow the user to log in. Specify when the user is warned about concurrent sessions:
 - Select **Always** to notify users each time they log in when they already have another active session
 - Select **If the maximum session has been exceeded** to display the warning message only when the user's maximum session count has been met.
5. **To specify the maximum number of concurrent sessions:**
 1. Select **Users > User Realms > *RealmName* > Authentication Policy > Limits**.
 2. Specify the number of sessions permitted for users in the **Maximum number of sessions per user text** box.
 3. Click **Save Changes**.



If you do not select the Enable multiple user sessions check box, only one session per user is allowed regardless of the value you specify in the **Maximum number of sessions per user** text box.

IF-MAP Federation Overview

You can configure a Ivanti Policy Secure device to store user session information for other Ivanti Policy Secure and Ivanti Connect Secure devices. Federation allows users to authenticate to a single Ivanti Connect Secure or Ivanti Policy Secure, and then access resources that are protected by any number of Ivanti firewall devices known as Infranet Enforcers that are controlled by different Infranet Controllers. Federation enhances network performance. If a user is required to log in to multiple Ivanti Connect Secure or Ivanti Policy Secure devices during the course of a day to access different resources, each device must perform authentication and Host-Checking, often with periodic Host Checker updates throughout the day. The overhead can lead to decreased performance not only on the devices, but also on the network and the endpoint. Imported IF-MAP sessions eliminate redundant logins and Host Checks.

Federation on the device uses the standard IF-MAP (Interface for Metadata Access Point) protocol to share session information and other data between connected devices over distributed networks. IF-MAP is a protocol defined by the Trusted Network Connect Working Group (TNC-WG) as a standard interface between different network elements and devices. Federation is accomplished using an IF-MAP server and IF-MAP clients.

It is important as an administrator to understand the fundamental underlying communication method for data transmission in a Federation network over IF-MAP. Policies that you configure on the device permit this communication.

In a federated network, the IF-MAP server functions as the repository, or data store for IF-MAP clients to use for publishing information regarding activity on the network. For example, IF-MAP clients can publish information about sessions on the network, and Juniper Networks IDP devices can communicate information about potential threats to the IF-MAP client for publishing. IF-MAP clients can search for information about sessions or threats, and an IF-MAP client can establish a subscription so the IF-MAP server notifies the client when other clients publish new or changed information. In addition, IF-MAP clients can purge data that is no longer valid. All transactions are initiated by the IF-MAP client.

IF-MAP Federation is available on all Ivanti Connect Secure devices with version 6.4 or greater. No licensing is required.

1. The endpoint authenticates through the IF-MAP client (Ivanti Connect Secure). The IF-MAP client publishes session information to the IF-MAP server.
2. The endpoint attempts to access protected resources that are behind the Infranet Enforcer.
3. The Infranet Enforcer notifies the Infranet Controller (also an IF-MAP client). The IF-MAP client searches for session information on the IF-MAP server.
4. The Infranet Controller subscribes to session information about the endpoint's IP address.
5. The Infranet Controller notifies the Infranet Enforcer that session information exists for the IP address attempting to access resources, and the Infranet Enforcer provisions an auth table entry.
6. Access is granted to the protected resources. If any session information about the user changes, the authenticating IF-MAP client publishes the new information. Having subscribed to the user's session information, the Infranet Controller will be aware of any changes and provision access in accordance with the changed session information.

For details about configuring the system to work in an IF-MAP Federated network with the Infranet Controller, see IF-MAP Feature Guide.

IF-MAP Federation Workflow

Configuring an IF-MAP Federated network requires coordination between administrators of the different devices that will be in the federated network.

This document describes IF-MAP deployments that include only Ivanti devices: Infranet Controllers, Ivanti Connect Secure devices, Infranet Enforcer firewalls, and Juniper Networks IDP. For implementations that incorporate third-party components, contact Ivanti Technical Support.

The mix of devices in the federated network is important when planning the network. Will your network consist of only Infranet Controllers, or will you incorporate Ivanti Connect Secure? Do the devices in your network have similar role mapping policies, or is each device different?

Determine and understand your goals for the federated network. The big picture guides your implementation as it becomes more complex. Ivanti recommends that you begin slowly. For example, start with a single role on each device, and then build the network incrementally.

In the simplest model, you can use the default policies. Using this model, you can quickly establish a federated network, and session information will automatically be shared among distributed devices in the network. This simple model should be adequate for most implementations in which the devices in the federated network have identical or very similar role mapping policies.

If your configuration requires more complex policies, you will need to perform a number of tasks to achieve your intended results. The following guidelines will help you plan your workflow:

- Ensure that communications between IF-MAP servers and IF-MAP clients is established
- Determine the resources that will be shared among the different devices
- Define who can access specific resources
- Distribute resources and users into roles
- Establish a naming convention that is shared and implemented between all administrators and devices
- Create Session-Export and Session-Import policies that reflect the role designations that you have configured on the devices

IF-MAP Federation Details

You can configure the system as an IF-MAP client for an IF-MAP server. You configure an Infranet Controller as an IF-MAP server. Any endpoint sessions with an IP address created on an IF-MAP server are automatically published to that IF-MAP server.

You can create source IP policies for endpoints that authenticate to the device to permit access to resources behind Infranet Enforcers (ScreenOS Enforcers and Ivanti Policy Secure s). Session-Export policies that you configure on the IF-MAP clients allow the clients to publish endpoint user data to the IF-MAP server. Devices that are IF-MAP clients can subscribe to the information on an IF-MAP server.

When a user accesses the device that is configured as an IF-MAP client, the client publishes basic session information, including the IP address, username and roles, to the IF-MAP server. The server stores the information as metadata. Other IF-MAP clients in the network can poll the server for metadata when session information is needed as a result of an endpoint attempting to access protected resources behind an Infranet Enforcer.

When an authenticated user from the device that is configured as an IF-MAP client attempts to access resources that are protected by an Infranet Enforcer for an Infranet Controller that is also configured as an IF-MAP client, the Infranet Controller automatically provisions an auth table entry for the user on the Infranet Enforcer to allow access without requiring the user to authenticate to the Infranet Controller.

The Infranet Enforcer as an IF-MAP client subscribes to session information and other data for the endpoint based on the originating IP address. The authenticating device (the original IF-MAP client) publishes any changes in session parameters to the IF-MAP server. Since the Infranet Controller that is protecting the accessed resources subscribes to the metadata on the Federation server, session information is always current.

The Infranet Enforcer allows or denies traffic based on the resource access policies that are configured on the Infranet Controller to which it is connected.

You configure server settings on the Infranet Controller that will be the IF-MAP server. You configure client settings on each of the Ivanti Connect Secure and Infranet Controller devices and that will be connected in the network.

In addition to the server and client settings, you configure Session-Export policies on Ivanti Connect Secure and Infranet Controllers that are IF-MAP clients You configure and Session-Import policies on Infranet Controller IF-MAP clients that are connected to Infranet Enforcers.


IF-MAP allows servers and clients to publish, search, poll, and subscribe to data within a network of IF-MAP servers and clients. All of the data from Ivanti Connect Secure in the network that is published to the IF-MAP server uses the IF-MAP protocol. Session-Export and Session-Import policies that you configure on Ivanti Connect Secure and Infranet Controller allow the devices to utilize the IF-MAP protocol to share session information.

Session-Export policies specify how to translate an endpoint's session on Ivanti Connect Secure or Ivanti Connect Secure into IF-MAP data. To translate session information into IF-MAP data, you enter detailed user information. Ivanti Connect Secure evaluates the Export policies to determine a session's IF-MAP roles, capabilities, identities, and device attributes and publishes the data to the IF-MAP server.

The Session-Import policies that you configure on Ivanti Policy Secure specify how the device should derive a username and a set of roles based on IF-MAP data that it receives from the IF-MAP server from other Ivanti Connect Secure devices. Import policies are similar to Role Mapping policies on a realm. You must be precise when configuring Export and Import policies, otherwise roles cannot be assigned properly.

The following figure depicts a scenario in which there are two Infranet Controllers configured as IF-MAP clients, one Ivanti Connect Secure device configured as an IF-MAP client, and another Infranet Controller configured as the IF-MAP server. Endpoints that authenticate through any of the IF-MAP clients can access protected resources behind the enforcement point attached to Infranet Controller 1.

The following figure depicts the IF-MAP Federated Network Topology:



The interaction between the endpoints, the clients and the server are as follows:

The interaction between the endpoints, the clients and the server are as follows:

- An endpoint authenticates through Ivanti Connect Secure depicted in the figure and starts VPN Tunneling or Pulse client.

- Ivanti Connect Secure provisions an IP address for the endpoint to use on the internal network. Once the endpoint's IP address on the internal network is known, Ivanti Connect Secure derives IF-MAP data from the endpoint's session.
- The Connect Secure IF-MAP client publishes the session information as IF-MAP data to the IF-MAP server using Session-Export policies.
- When the user attempts to access resources behind the enforcement point, access is blocked since the Infranet Enforcer has no information about the endpoint. The Infranet Enforcer sends out a dynamic discovery message that includes the endpoint's source IP address.
- Infranet Controller 1 uses the IP address to retrieve session data from the IF-MAP server.
- The Infranet Controller uses Session-Import policies to retrieve session data from the IF-MAP server.

The endpoint authenticating to Ivanti Connect Secure must be running VPN Tunneling.

Imported user sessions do not count against the maximum user count for either platform, as each user is counted on the Ivanti Connect Secure device from which they authenticated.

For details on configuring an IF-MAP Federation network, see *IF-MAP Feature Guide*.

IF-MAP Logging

IF-MAP related events are logged on both the IF-MAP server and the IF-MAP client.

Task Summary: Configuring IF-MAP Federation

The tasks listed in this topic are performed by a Ivanti Policy Secure administrator, in conjunction with an administrator for Ivanti Connect Secure. On Ivanti Connect Secure, you configure Session-Export policies and you configure IF-MAP client settings. For details on configuring an IF-MAP Federated network, see the *IF-MAP Feature Guide*.

To use IF-MAP Federation, perform the following tasks on Ivanti Policy Secure and Ivanti Connect Secure:

1. Enable dynamic auth table provisioning on any connected Infranet Enforcers that you want to use with Federation.
2. On Ivanti Policy Secure, configure IF-MAP server settings to permit the server to communicate with IF-MAP clients.
3. Configure IF-MAP client settings to permit clients to communicate with the IF-MAP server.

4. On Ivanti Policy Secure and Ivanti Connect Secure, coordinate Session-Import policies, Session-Export policies, roles, and resource access policies between all of the clients in the Federated network.
5. Configure Session-Export policies on Connect Secure to define how sessions are translated into IF-MAP data.
6. Configure Session-Import policies on Connect Secure that correspond with Export policies to translate IF-MAP data into roles.
7. On Policy Secure, configure Source IP policies for Connect Secure users who will use Source IP to access the network.

Configuring IF-MAP Server Settings

You must add all IF-MAP clients to the IF-MAP server. To add clients, you must specify the IP address and the security mechanism and credentials for the client.

For details on configuring an IF-MAP server, see *IF-MAP Feature Guide*.

Configuring the IF-MAP Federation Client

You must identify the IF-MAP server to each IF-MAP client. To add the server, you specify the IF-MAP URL of the server and how to authenticate to the server. Match the URL and security settings to equal those on the IF-MAP server(s) to which the IF-MAP client will connect.

To configure IF-MAP client settings on the devices that will be IF-MAP clients:

1. From the admin console select System > **IF-MAP Federation** > **Overview**.
2. Select the **Enable IF-MAP Client** check box.
3. Type the **Server URL** for IF-MAP Web service on the IF-MAP server. Append the server URL with **/dana-ws/soap/dsifmap** for all Ivanti IF-MAP servers.
4. Select the client Authentication method: **Basic** or **Certificate**.
 - If you select **Basic**, enter a **Username** and **Password**. This is the same as the information that was entered on the IF-MAP server.
 - If you select **Certificate**, select the **Device Certificate** to use.

- Ensure that the certificate of the CA that signed the IF-MAP server certificate is added from the **System > Configuration > Certificates > Trusted Server CA** page. The IF-MAP client validates the IF-MAP server certificate: if validation fails, the connection fails. Ensure that the hostname in the IF-MAP URL on the client machine matches the hostname of the server certificate on the IFMAP server, and that the CA that signed the server certificate is configured as a trusted server CA on the IF-MAP client.
5. Click **Save Changes**.

IF-MAP Federated Network Timing Considerations

It is important that the time on all IF-MAP servers is correct, as timeout issues are critical to ensure that IF-MAP provides complete and timely information. The IF-MAP Federation is designed to fail secure. If any component in the network does not receive timely information, the IF-MAP metadata will be purged from the data stores.

The components are designed to fail-secure. If complete and timely information cannot be provided, a user's session will be deleted. For example, if the chain of connections between an IF-MAP client that publishes a session and a client that grants access to a resource breaks, the client that granted access will remove the session. The fail-secure time limit is three minutes.

The timeout limit for IF-MAP is three minutes and applies to the following events:

- An IF-MAP server (or cluster) loses contact with one of its IF-MAP clients
- An IF-MAP server (cluster) loses contact with one of the other IF-MAP server (clusters) in the IF-MAP federation
- A Ivanti IF-MAP client loses contact with its IF-MAP server (cluster) for too long

Session-Export and Session-Import Policies

You configure Session-Export policies on all of the Connect Secure and Infranet Controller devices in the Federated network that are IF-MAP clients. These policies allow IF-MAP clients to translate outgoing session information into IF-MAP data and incoming IF-MAP data into session information. These translations enable sessions to be shared between Connect Secure and Infranet Controller devices even if the devices sharing sessions have different role configurations.

To accurately configure Session-Export and Session-Import policies you need a minimal understanding of IF-MAP identifiers and IF-MAP metadata. An identifier is a unique value required for all metadata operations. Each instance of metadata is associated with an identifier. Examples of identifiers include access-request, identity, IP address, and MAC address. Examples of metadata include capability, role, and device-attribute.


IF-MAP recognizes two metadata types that are similar to roles on Connect Secure: IF-MAP roles and IF-MAP capabilities. An IF-MAP role is an attribute assigned to a user in the organization. When IF-MAP metadata is published to the IF-MAP server, this information could be one way to identify individuals on the network. This is somewhat different from the concept of roles. An IF-MAP capability is closer to the concept of a role. An IF-MAP capability is a collection of privileges assigned as a result of an access request. This is more analogous to a role since they are derived through role mapping in an authentication realm.

The data that is published to the IF-MAP server about a user session is derived by applying the Session-Export policies to the user session. The Session-Import policies are applied to the data from the IF-MAP server to assign a set of roles to the user.

When an endpoint attempts to access protected resources associated with Connect Secure, the device queries the IF-MAP server for data. The Infranet Controller uses Session-Import policies to derive roles and a username from the IF-MAP data. For example, you could configure a Session-Import policy that looks for a specific Host Checker policy (you specify the Host Checker policy in the Session-Import policy). If the Infranet Controller finds a match (in this case the Host Checker device attribute), the user can be assigned a role specified in the Session-Import policy.

All of the administrators who are configuring devices in the IF-MAP Federated network must agree on a set of capabilities, roles and device attributes and their meanings to be used with IF-MAP. Then, each administrator configures their device to map between local sessions and IF-MAP data. Figure 5 illustrates a coordinated IF-MAP Federated network configuration with policies that permit an example user to access protected resources.

The following figure depicts Session-Import and Session-Export Policies:



To further your understanding of Session-Import and Session-Export policies, please note the following Ivanti IF-MAP conventions:

To further your understanding of Session-Import and Session-Export policies, please note the following Ivanti IF-MAP conventions:

- The system maps to the identical IF-MAP username.
- A role is paired with an IF-MAP capability.
- Capabilities can have the same name as the roles they are paired with, or a different name.
- When different IF-MAP clients have different but equivalent role names (e.g. Legal and Law, both referring to members of the corporate legal department) a single IF-MAP capability must be chosen.
- Not every role needs to be paired with an IF-MAP capability: roles can be local to Connect Secure.
- After you decide on pairings between IF-MAP capabilities and the roles, you create a session export policy for each pairing. On an Infranet Controller that controls Infranet Enforcers, you create a session import policy.
- The only parameters for the policies are the roles and the IF-MAP capability; everything else is fixed.

Default Session-Export and Session-Import Policy Action

By default, Session-Import and Session-Export IF-MAP policies are configured to allow IF-MAP capabilities (the equivalent of roles) to be published to the IF-MAP server and retrieved from the IF-MAP server, provided there are matching roles on each IF-MAP client. You can open new Session-Import and Session-Export policies on each device, and then name and close the policies. Any matching roles that the IF-MAP clients in the federated network have can be used to access resources.

Advanced Session-Export and Session-Import Policies

By default, advanced policy actions are not visible unless you click the advanced options links on the Session-Export and Session-Import policy pages. In default mode, you configure Session-Export and Session-Import policies using IF-MAP capabilities and roles.

Device attributes, IF-MAP roles and identities can be accessed through the advanced options links. IF-MAP capabilities and Connect Secure roles should provide the functionality that most IF-MAP Federation requires.

Configuring Session-Export Policies

Session-Export policies determine how users are identified on the IF-MAP server when their session is published via IF-MAP: the policy sets the IF-MAP identifiers. You define attributes for users that will be used to determine role matching on different Infranet Controllers. For example, you might configure a Session-Export policy to specify that any users that belong to the "engineering" role should be identified with the "engineering" IF-MAP capability on the IF-MAP server. That identity will be included in the session information to which other IF-MAP clients subscribe. You configure corresponding Session-Import Policies on Infranet Controllers to identify which roles the user should be assigned.

You configure Session-Export policies based on Infranet Controller or Connect Secure roles, and users belonging to those roles can access resources on an Infranet Enforcer only if the role can be successfully matched with a role on the target Infranet Controller. You configure Session-Export policies on all Infranet Controller and Connect Secure devices for which you have users that will be allowed to access resources behind an Infranet Enforcer in the network.

When a user for whom Session-Export policies has been configured successfully authenticates to the network, the Session-Export policies are used to translate the user session into IF-MAP data which is then sent to the IF-MAP server. When the user attempts to access a resource that is protected by an Infranet Enforcer, the target Infranet Controller then attempts to translate the IF-MAP data for the user into a username and roles using the Session-Import policies that are configured on the second Infranet Controller device.

Administrative Domains in Session-Export Policies

In a Layer 2 environment, session information on the IF-MAP server includes a MAC address. If an export policy specifies an Administrative Domain, the domain is associated with the MAC address published to the IF-MAP server (the administrative domain is also associated with the identity published to the IF-MAP server).

A DHCP server assigns an IP address to the endpoint after authentication. An IF-MAP enabled DHCP server publishes an ip-mac link to IF-MAP, associating the endpoint's IP address with its IF-MAP session information.

Including administrative domains in MAC addresses allows the ip-mac link to be created based on the administrative domain.

If your IF-MAP Federated network spans different administrative domains, you should configure separate Session-Export policies for each domain to prevent MAC address spoofing. Each administrative domain should have an associated DHCP server and unique Session-Export policies.

Other aspects of the Session-Export policies within the IF-MAP Federated network can overlap.

To configure a Session-Export policy:

1. From the admin console select **System > IF-MAP > Session-Export Policies**.
2. Click **New** to create a new policy.
3. Type a **Policy Name**, and optionally a **Description**.
4. Optionally, add **Available Roles** to the **Selected Roles** column to determine the roles for which this policy should apply. If you do not add any roles, the policy applies to all sessions. However, if you have non-interactive devices such as printers that do not need access, you may want to manually add roles and exclude those roles with non-interactive devices.
5. Under Policy Actions, Select **Set IF-MAP Capabilities** and choose the applicable roles.
 - **Copy matching roles** - Selecting this action copies all of the user roles that match the roles specified in the Roles section of this policy into the IF-MAP capabilities data.
 - **Copy ALL roles** - Selecting this action copies all of the roles from the user session to the IF-MAP capabilities data.
 - **Set capabilities specified below** - Enter capabilities, one per line.
6. Select **Stop processing policies when this policy matches** to specify that when this policy is matched, no more Session-Export policies should be applied.
7. Click **Save Changes**, or continue to configure Advanced Actions.

To configure advanced actions (generally not required for Infranet Controller and Connect Secure IF-MAP Federation):

1. Click the **View Advanced Actions** link. Additional options appear on the page.
2. **Set IF-MAP Identity** - If this action is chosen, enter the **Identity** and select an **Identity Type** from the menu. Identity is normally specified as <NAME>, which assigns the user's login name. Any combination of literal text and context variables may be specified. If you choose **other** for Identity Type, enter a unique Identity Type in the **Other** text box.
3. Optionally type the **Administrative Domain** for the Session-Export policy. This optional field is applied to identity and MAC address data. One example for using this field is in a large network environment with several domains in which a username could be duplicated. By entering the domain, you ensure that the correct user is identified.
4. **Set IF-MAP Roles** - If this action is selected, select the applicable roles.
 - **Copy matching roles** - Selecting this action copies all of the user roles that match the roles specified in the Roles section of this policy into the IF-MAP capabilities data.
 - **Copy ALL roles** - Selecting this action copies all of the roles from the user session to the IF-MAP capabilities data.
 - **Set roles specified below** - Enter roles, one per line.
5. **Set IF-MAP Device Attributes** - Device attributes represent a passed Host Checker policy on the Infranet Controller or Connect Secure.
 - **Copy Host Checker policy names** - The name of each Host Checker policy that passed for the session is copied to a device attribute.
 - **Set device attributes specified below** - Type device attributes, one per line, into the text box.
6. Click **Save Changes** to save this advanced Session-Export policy.

You must create corresponding Session-Import policies that allow IF-MAP client Infranet Controllers that are connected to an Infranet Enforcer in front of protected resources to collect IF-MAP data from the IF-MAP server.

Session-Import Policies

The Session-Export policies that you create allow IF-MAP data that represents a session to be stored on the IF-MAP server. Session-Import policies specify how the Infranet Controller derives a set of roles and a username from the IF-MAP data in the IF-MAP server. Session-Import policies establish rules for importing user sessions from Connect Secure. Import policies allow you to match authenticated users with corresponding roles on the target device. For example, you might configure an Import policy to specify that when IF-MAP data for a session includes the "Contractor" capability, the imported session should have the "limited" role. Session-Import policies allow the Infranet Controller to properly assign roles based on information that the IF-MAP server provides.

You configure Session-Import policies on IF-MAP client IVEs that are connected to an Infranet Enforcer in front of protected resources. For information about configuring Session-Import policies, see IF-MAP Feature Guide.

Troubleshooting the IF-MAP Federated Network

Diagnostic tools on the Infranet Controller and Connect Secure can assist you with troubleshooting a federated network.

IF-MAP Client User Messages - On the IF-MAP client, logs information that is published and removed from the IF-MAP server.

- Enable **IF-MAP Client User Messages** from **Log/Monitoring > User Access > Settings** on the Infranet Controller and Connect Secure IF-MAP client.

IF-MAP Server Trace - On the IF-MAP server, logs the XML for all IF-MAP requests and responses.

- Enable **IF-MAP Server Trace** from **Log/Monitoring > Events > Settings** on the IF-MAP server.

IF-MAP Server Trace should only be enabled for troubleshooting purposes, as running this diagnostic incurs a large performance impact.

Viewing Active Users on the IF-MAP Client

On an IF-MAP client, you can view all of the sessions from other Infranet Controllers or Connect Secure devices that currently access the client (the imported sessions). Session information that can be viewed includes the username, roles, the user's endpoint IP address, and the IP address of the Infranet Controller or Connect Secure device that authenticated the user. You can select and remove sessions either temporarily or permanently. A temporarily removed session can be restored in response to a request for continued access. A permanently removed session cannot be restored.

To view, de-activate, or activate current sessions on an IF-MAP client:

1. Select **System > IF-MAP > Active Users from the IF-MAP** client admin console.
2. Select **Imported** or **Exported**.
3. Select **Activate** or **De-activate**.

Trusted Server List

The system uses two mechanisms to install and launch client software from a web browser:

- ActiveX controls (available only for Windows/IE)
- Java applets

With both mechanisms, the user is prompted to trust ActiveX controls and Java applets they have not run before. Inherent problems with these types of mechanisms are:

- When the user trusts an ActiveX control that control is trusted forever.
- When trusting a Java applet, users are trusting all code that is signed by the exact same code signing certificate.

To address the above, administrators can create a text file (called a whitelist) that contains a list of trusted devices, fully qualified domain names or IP addresses, one per line. Administrators can configure two types of whitelists:

- **Admin whitelist** - The admin whitelist file can be modified only by the endpoint administrator. The administrator must use SMS or other mechanism to copy the admin whitelist file to the end-user's system. Admin whitelist files are located in:

%ProgramFiles%\Pulse Secure\Whitelist.txt (Windows)

/usr/local/pulsesecure/whitelist.txt (Macintosh and Linux)

- **User whitelist** - Users can themselves make the decision to trust a device. When the user makes a decision to trust device, the device gets added to the user whitelist. User whitelist files are located in:

%AppData%\Pulse Secure\Whitelist.txt (Windows)

/~/Library/Application Support/Pulse Secure/whitelist.txt (Macintosh)

/~/.pulse_secure/whitelist.txt (Linux)



The trusted server list feature is for applications launched from a browser window. It does not apply to applications launched from the command-line or other means.

Administrator and User Configuration

The following is a snippet of a whitelist file:

```
qa.pulsesecure.netdev1.pulsesecure.net66.129.224.48
```



Whitelist files are not deleted when the software is removed.

There are two modes of enforcement:

- **Allow Admin List Only** - When software launches from the device that is not in the administrator whitelist, the launch fails and the user receives the error message "You are not allowed to launch software downloaded from <server>. Contact your system administrator for assistance." If the device is in the administrator whitelist, the launch proceeds as requested.
- **Prompt** - When software launches from the device that is not in the administrator whitelist or the user whitelist, the user is prompted if they want to launch the software with the message "Do you want to download, install and/or execute software from the following server". If the user declines, the launch fails. If the user accepts, the launch proceeds. The user also has the option to automatically add the device to the user whitelist file by selecting one of the following options from the message window:
 - **Always** - Add the server to the user whitelist file and download, install or launch the software
 - **Yes** - Download, install or launch the software but don't add the server to the user whitelist file

- **No** - Do not download, install or launch software and don't add the server to the user whitelist file

If the first line of the whitelist file contains "AllowAdminListOnly" (case insensitive) then Allow Admin List Only enforcement mode is used. Otherwise, prompt mode enforcement is used.

A snippet of a whitelist file using Allow Admin List Only enforcement is shown here:

AllowAdminListOnly qa.pulsesecure.net dev1.pulsesecure.net 66.129.224.48



Prompt enforcement is the default mode when you upgrade your software to the latest revision.

To add clusters to the whitelist file:

- For Active/Passive clusters, enter the VIP in the whitelist.
- For Active/Active clusters, enter the load balancer hostname in the whitelist.

White List Flow Chart

The following steps outline the process for determining whether to launch the software

1. If the URL of the page initiating the launch does not begin with https, abort the launch and notify the user.
2. Else if the admin whitelist exists,
 - If the origin site is listed in the whitelist, proceed with the launch.
 - If the origin site is not in the whitelist and the whitelist starts with "AllowAdminListOnly", abort the launch and notify the user.
3. Else if the user whitelist exists,
 - If the origin site is in the user whitelist, proceed with the launch.
4. Prompt the user if they trust the origin site.
5. If the user agrees to trust the origin:
 - If they select **Always**, then add the server to user whitelist file.
 - Proceed with the launch.
6. Abort the launch.

User Roles

User Roles Overview

A user role is an entity that defines user session parameters (session settings and options), personalization settings (user interface customization and bookmarks), and enabled access features (Web, file, secure application manager, VPN tunneling, Secure Email, enterprise onboarding, Terminal Services, e-mail access, virtual desktops, HTML5 access, and Pulse client). A user role does not specify resource access control or other resource-based options for an individual request. For example, a user role may define whether or not a user can perform Web browsing. However, the individual Web resources that a user may access are defined by the Web resource policies that you configure separately.

The access management framework supports two types of user roles:

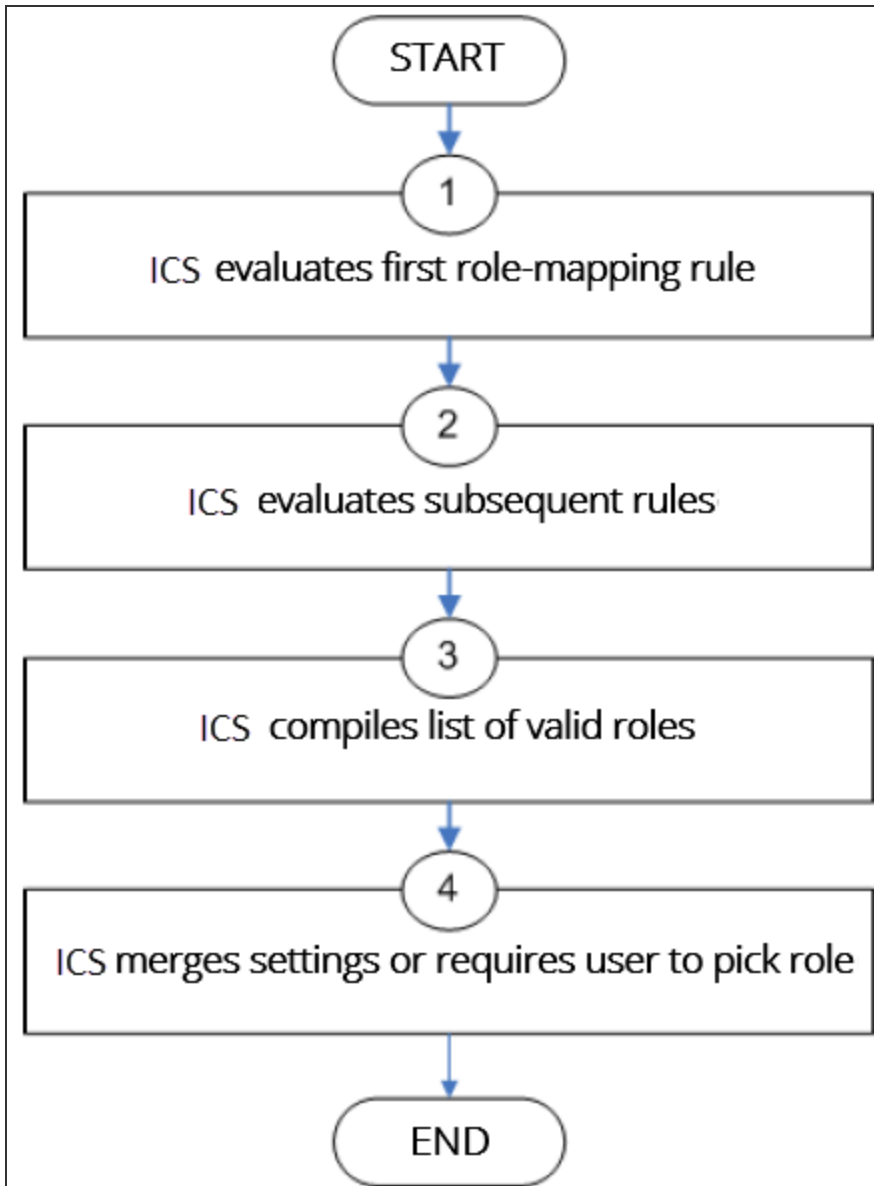
- **Administrators** - An administrator role specifies management functions and session properties for administrators who map to the role. You can customize an administrator role by selecting the feature sets and user roles that members of the administrator role are allowed to view and manage. You can create and configure administrator roles through the Delegated Admin Roles page. Click Administrators > Admin Roles in the admin console.
- **Users** - A user role is an entity that defines user session parameters, personalization settings, and enabled access features. You can customize a user role by enabling specific access features, defining Web, application, and session bookmarks, and configuring session settings for the enabled access features. You can create and configure user roles through the Roles page. Click **Users > User Roles** in the admin console.

User roles are an integral part of the access management framework, and therefore are available on all Ivanti Connect Secure products. However, you can only access features through a user role if you are licensed for the feature.

User Role Evaluation

The role mapping engine determines a user's session role, or combined permissions valid for a user session, as illustrated in the following figure. A detailed description of each step follows the diagram.

The following figure depicts the Security Checks Performed by Ivanti Connect Secure to Create a Session Role:



The system performs the following security checks to create a session role:

1. The system begins rule evaluation with the first rule on the Role Mapping tab of the authentication realm to which the user successfully signs in. During the evaluation, the system determines if the user meets the rule conditions. If so, then:
 - The system adds the corresponding roles to a list of "eligible roles" available to the user.

- The system considers whether or not the "stop on match" feature is configured. If so, then the engine jumps to step 5.
2. The system evaluates the next rule on the authentication realm's Role Mapping tab according to the process in Step 1 and repeats this process for each subsequent rule. When the system evaluates all role mapping rules, it compiles a comprehensive list of eligible roles.
 3. The system evaluates the definition for each role in the eligibility list to determine if the user complies with any role restrictions. The system then uses this information to compile a list of valid roles, whose requirements the user also meets.

If the list of valid roles contains only one role, then the system assigns the user to that role. Otherwise, the system continues the evaluation process.

4. The system evaluates the setting specified on the Role Mapping tab for users who are assigned to more than one role:
 - **Merge settings for all assigned roles** - If you choose this option, then the system performs a permissive merge of all the valid user roles to determine the overall (net) session role for a user session.
 - **User must select from among assigned roles** - If you choose this option, then the system presents a list of eligible roles to an authenticated user. The user must select a role from the list, and the assigns the user to that role for the duration of the user session.
 - **User must select the sets of merged roles assigned by each rule** - If you choose this option, the system presents a list of eligible rules to an authenticated user (that is, rules whose conditions the user has met). The user must select a rule from the list, and the system performs a permissive merge of all the roles that map to that rule.



If you use automatic (time-based) dynamic policy evaluation or you perform a manual policy evaluation, the system repeats the role evaluation process described in this section.

Permissive Merge Guidelines

A permissive merge is a merge of two or more roles that combines enabled features and settings following these guidelines:

- Any enabled access feature in one role takes precedence over the same feature set disabled in another role.

- In the case of Secure Application Manager, the system enables the version corresponding to the first role that enables this feature. Furthermore, the system merges the settings from all the roles that correspond to the selected version.



If you are using Pulse client, then Ivanti is always enabled as the default client.

- In the case of user interface options, the system applies the settings that correspond to the user's first role.
- In the case of session timeouts, the system applies the greatest value from all of the roles to the user's session.
- If more than one role enables the Roaming Session feature, the system merges the netmasks to formulate a greater netmask for the session.
- When merging two roles that a user is mapped to-one in which bookmarks open in a new window and one in which bookmarks open in the same window-the merged role opens bookmarks in the same window.
- When merging two roles in which the first role disables the browsing toolbar and the second role enables either the framed or standard toolbar, the merged role uses the settings from the second role and displays the specified browsing toolbar.
- The merged role uses the highest value listed for each HTTP Connection Timeout. Click **Users > User Roles > Select Role > Web > Options** then click **View** advanced options.
- Merging of conflicting VPN Route Precedence Options is discouraged. But if it is done, the order of precedence is Allow Local Subnet Access, then Tunnel Routes and then Endpoint Routes.

Configuration of User Roles

To create a user role:

1. In the admin console, choose **Users > User Roles**.
2. Click **New Role** and then enter a name and optionally a description. This name appears in the list of Roles on the Roles page.

Once you have created a role, you can click the role's name to begin configuring it using the instructions in the following sections.

When you delete a role, the personal bookmarks, SAM settings, and other settings may not be removed. Therefore, if you add a new role with the same name, any users added to that new role may acquire the old bookmarks and settings. In general, the system enforces referential integrity rules and does not allow you to delete any objects if they are referenced elsewhere.



For example, if a role is used in any of the realm's role mapping rules, then the system rejects the deletion of the role unless you modify or delete the mapping rules.

When you create individual user accounts, you must add the users through the appropriate authentication server (not the role). Or for instructions on how to create users on third-party servers, see the documentation that comes with that product.

Configuring General Role Options

Click Overview at the top of the General tab to edit a role's name and description, toggle session and user interface options on and off, and enable access features. When you enable an access feature, make sure to create corresponding resource policies.

To manage general role settings and options:

1. In the admin console, click **Users > User Roles > Role Name > General > Overview**.
2. Revise the name and description and then click **Save Changes (optional)**.
3. Under **Options**, select the role-specific options that you want to enable for the role.

The system uses default settings for newly created roles or when you do not select role-specific options.

Role-specific options include:

- **VLAN/Source IP** - Select this option to apply the role settings configured on the General > VLAN/Source IP page.
 - **Session Options** - Select this option to apply the role settings in the General > Session Options page to the role.
 - **UI Options** - Select this option to apply the role settings in the General > UI Options page to the role.
4. Under Access features, select the features you want to enable for the role. Options include:
 - **Web** - intermediate Web URLs through the Content Intermediation Engine.

- **Files (Windows)** - resource profile that controls access to resources on Windows server shares.
 - **Secure Application Manager (Windows version or Java version)** - provides secure, application-level remote access to enterprise servers from client applications.
 - **Terminal Services** - enable terminal emulation sessions on a Windows terminal server, Citrix NFuse server, or Citrix Metaframe server.
 - **VPN Tunneling** - provides secure, SSL-based network-level remote access to all enterprise application resources using the system.
 - **Secure Mail** - enables automatic synchronization with an Exchange server (ActiveSync) and e-mail encryption for iOS devices that have the Pulse client.
 - **Enterprise Onboarding** - allows users to securely access enterprise network resources with almost any device. Wi-Fi, VPN, certificate, and Secure Mail profiles can be defined for enterprise resources and downloaded to a device during onboarding, depending on the device type.
5. Click **Save Changes** to apply the settings to the role.

Role Restrictions

Click **Restrictions** at the top of the **General** tab to specify access management options for the role. The system considers these restrictions when determining whether or not to map a user to the role. The system does not map users to this role unless they meet the specified restrictions.

You may configure any number of access management options for the role. If a user does not conform to all of the restrictions, the system does not map the user to the role.

To specify access management options for the role:

1. In the admin console, click **Users > User Roles > Role Name > General > Restrictions**.
2. Click the tab corresponding to the option you want to configure for the role, and then configure it.

Specifying Role-Based Source IP Aliases

Click VLAN/Source IP at the top of the General to define role-based source IP aliases. If you want to direct traffic to specific sites based on roles, you can define a source IP alias for each role. You use these aliases to configure virtual ports you define for the internal interface source IP address. A back-end device can then direct end user traffic based on these aliases, as long as you configure the back-end device, such as a firewall, to expect the aliases in place of the internal interface source IP address. This capability enables you to direct various end users to defined sites based on their roles, even though all of the end user traffic has the same internal interface source IP address.



You must define virtual ports to take advantage of the role-based source IP aliases.

To specify a source IP alias for the role:

1. In the admin console, click **Users > User Roles > Role Name > General General > VLAN/Source IP**.
2. Select the **VLAN** you want to use from the VLAN list, if you have defined VLAN ports on your system.

If you have not defined VLAN ports, the option defaults to the Internal Port IP address.

3. Select a source **IP address** from the list.
4. Click **Save Changes** to apply the settings to the role.



If an end user is mapped to multiple roles and the system merges roles, the system associates the source IP address configured for the first role in the list with the merged role.

You can specify the same source IP address for multiple roles. You cannot specify multiple source IP addresses for one role.

Specifying Role Session Options

Use the Session tab to specify session time limits, roaming capabilities, session and password persistency, request follow-through options, and idle timeout application activity. Select the Session Options check box on the Overview tab to enable these settings for the role.

To configure general session options:

1. In the admin GUI, click **User > User Roles > RoleName > General > Session Options**.
2. Configure session options, as described in [Table](#).

3. Save the configuration.

The following table lists Session Options

Option	Guidelines
Session lifetime	<p>For Idle Timeout, specify the number of minutes a non-administrative user session may remain idle before ending. The minimum is five minutes. The default idle session limit is 10 minutes, which means that if a user's session is inactive for 10 minutes, the system ends the user session and logs the event in the system log (unless you enable session timeout warnings described later).</p> <p>For Max. Session Length, specify the number of minutes an active non-administrative user session may remain open before ending. The minimum is six minutes. The default time limit for a user session is 60 minutes, after which the system ends the user session and logs the event in the system log. During an end user session, prior to the expiration of the maximum session length, the system prompts the user to reenter authentication credentials, which avoids the problem of terminating the user session without warning.</p> <p>For Reminder Time, specify when the system should prompt non-administrative users, warning them of an impending session or idle timeout. Specify the number of minutes before the timeout is reached.</p> <p>Optionally, select Use Session/Idle timeout values sent by the primary Radius authentication Server to override the idle timeout and session length specified above. If the received values are below the minimums (5 minutes for the idle timeout and 6 minutes for the session length), the minimum values are used.</p> <p>Optionally, select Enable Session Extension to allow users to extend the session beyond the maximum session length. If this feature is enabled, users will be reauthenticated and extend their current session without interruption.</p> <p>We recommend the difference between Idle Timeout and Reminder Time be greater than two minutes. This ensures that the reminder pop-up window appears at the correct time.</p>

Option	Guidelines
Enable session timeout warning	<p>Enable to notify non-administrative users when they are about to reach a session or idle timeout limit.</p> <p>These warnings prompt users to take the appropriate action when they are close to exceeding their session limits or idle timeouts, helping them save any in-progress form data that would otherwise be lost. Users approaching the idle timeout limit are prompted to reactivate their session. Users approaching the session time limit are prompted to save data.</p> <p>Optionally, select Display sign-in page on max session time out to display a new browser sign-in page to the end user when their session times out. This option only appears when you choose to enable the session timeout warning.</p> <p>If you do not select the Enable session timeout warning option, the system only displays expiration messages to users. It does not give them the option to extend their sessions. Instead, users need to access the sign-in page and authenticate into a new session.</p> <p>The Enable session timeout warning option only applies to expiration messages displayed by the end user's browser, not by other clients such as PSAM or VPN Tunneling.</p>
Roaming session	<p>Select one of the following options:</p> <p>Enabled - Enables unlimited roaming sessions. An unlimited roaming session allows mobile users (laptop users) with dynamic IP addresses to sign in to the device from one location and continue working from another. Disable this feature to prevent users from accessing a previously established session from a new source IP address. This helps protect against an attack spoofing a user's session, provided the hacker was able to obtain a valid user's session cookie. If you enable unlimited session roaming, a session is maintained within an IPv4 network, within an IPv6 network, or from IPv4 to IPv6 and vice versa.</p> <p>Limit to subnet - Limits the roaming session to the local subnet, specified by netmask for IPv4 subnets and prefix length for IPv6 subnets. Users may sign in from one IP address and continue using their sessions with another IP address as long as the new IP address is within the same subnet. If you configure limited session roaming, you can specify IPv4 or IPv6 subnets within which the session is maintained. However, with limited session roaming, you cannot allow sessions to roam from IPv4 to IPv6 networks, or vice versa.</p> <p>Disabled - Disables roaming user sessions for users mapped to this role. Users who sign in from one IP address may not continue an active session from another IP address; user sessions are tied to the initial source IP address.</p>

Option	Guidelines
Persistent session	<p>By default, the session cookie is flushed from the browser's memory when the browser is closed. The session length is determined by both the idle timeout value and maximum session length value that you specify for the role. The session does not terminate when a user closes the browser; a session only terminates when a user signs out of the device.</p> <p>Enable the persistent session option to write the session cookie to the client hard disk so that the user's credentials are saved for the duration of the session. Assume persistent session is enabled and a user starts a VPN Tunneling session from a browser. Later, the user quits the browser application. The next time the user opens a new browser window and logs in to the same device, the user is not prompted to enter his or her credentials again.</p> <p>(Macintosh only) Persistent session applies only for browser login as stated earlier. If you start VPN Tunneling from the standalone launcher (by opening NetworkConnect.dmg) and later open a new browser and log in to that same device, you are prompted to reenter your credentials.</p> <p>If you enable the Persistent session option and a user closes the browser window without signing out, any user can open another instance of the same browser to access the device without submitting valid credentials, posing a potential security risk. We recommend that you enable this feature only for roles whose members need access to applications that require system credentials and that you make sure these users understand the importance of signing out of the device when they are finished.</p>

Option	Guidelines
Remove Browser Session Cookie	<p>Enable to remove the session cookie and logs users out of their Web session once the client component launches, enhancing security for your VPN Tunneling, PSAM and Ivanti session.</p> <p>Disable to retain the session cookie and keep users logged in to their Web session once the client component starts.</p> <p>Because browser cookies are plain text files, they are susceptible to malicious attacks. The Remove Browser Session Cookie option removes the session cookie, making your VPN Tunneling, PSAM and Pulse client sessions more secure. When enabled, users are logged out of their Web session once the client component (for example, Pulse client) launches. Users are logged out of their Web session regardless of whether the client component launches successfully or not. If the client component does not successfully launch, users can restart their Web session and try launching their client component again. This option also prevents any client component from launching a browser through the client. The Remove Browser Session Cookie removes only the session cookie. It does not remove non-system cookies or other any other cookie.</p>
HTTP Only Device Cookie	<p>Enable to set a HTTP only cookie along with DSID.</p> <p>This cookie cannot be read with the help of scripts and protects against XSS attacks and cookie stealing. This cookie along with DSID will be used to restore a user session.</p>
Persistent password caching	<p>Enable to allow cached passwords to persist across sessions for a role.</p> <p>The system supports Windows NT LAN Manager (NTLM) authentication protocol and HTTP Basic Authentication and supports servers that are set up to accept NTLM sign-in. The system caches NTLM and HTTP Basic Authentication passwords provided by users so that the users are not repeatedly prompted to enter the same credentials used to sign in to the server or another resource in the NT domain. By default, the system flushes cached passwords when a user signs out. A user can delete cached passwords through the Advanced Preferences page. After the end user logs in to the device, click Preferences and then click the Advanced tab.</p>
Browser request follow-through	<p>Enable to allow the system to complete a user request made after an expired user session after the user reauthenticates.</p>
Idle timeout application activity	<p>Enable to ignore activities initiated by Web applications (such as polling for e-mails) when determining whether a session is active. If you disable this option, periodic pinging or other application activity may prevent an idle timeout.</p>

Option	Guidelines
Upload Logs	Enable to allow the user to transmit (upload) client logs to the system. Use the System > Log/Monitoring > Client Logs > Settings page to completely enable client-side logs for the user.

Customizing the Welcome Page

Click **UI Options** at the top of the General tab to specify customized settings for the welcome page and the browsing toolbar for users mapped to this role. The welcome page (or home page) is the Web interface presented to authenticated users.

Click **Overview** at the top of the General tab, and then select the **UI Options** check box to enable custom settings for the role; otherwise, the system uses the default settings.

Personalization settings include the sign-in page, page header, page footer, and whether or not to display the browsing toolbar. If the user maps to more than one role, then the system displays the user interface corresponding to the first role to which a user is mapped.

To customize the welcome page for role users:

1. Click **Users > User Roles > RoleName > General > UI Options**.
2. Under Header, specify a custom logo and alternate background color for the header area of the welcome page (optional):
 - Click **Browse** and locate your custom image file. The new logo appears in the Current appearance box only after you save your changes.



You can only specify a JPEG or GIF file for a custom logo image. Other graphics formats are not displayed properly in the JSAM status window on some OS platforms.

- Type the hexadecimal number for the background color or click the **Color Palette** icon and pick the desired color. The Current appearance box updates immediately.
3. Under Sub-headers, select new background and text colors (optional):
 - Type the hexadecimal number for the Background color or click the **Color Palette** icon and pick the desired color. The Current appearance box updates immediately.
 - Type the hexadecimal number for the Text color or click the **Color Palette** icon and pick the desired color. The Current appearance box updates immediately.

4. Under **Start page**, specify the start page that you want users to see after they sign in and when they click the Home icon on the toolbar:
 - **Bookmarks page** - Select this option to display the standard Bookmarks page.
 - **Custom page** - Select this option to display a custom start page and then specify the URL to the page. The system rewrites the URL and creates an access control rule to allow users access to the URL. (Note that users can also enter the custom URL in the Browse field on the toolbar.) The system evaluates the access control rule after all other policies, which means another policy could deny access to the URL.
 - **Also allow access to directories below this url** - Select this option to allow users access to subdirectories of the custom-page URL. For example, if you specify `http://www.domain.com/`, users can also access `http://www.domain.com/dept/`.

5. Under **Bookmarks Panel Arrangement**, arrange the panels as you want to display them on the user's bookmarks page:

1. To select the name of a panel, click in the **Left Column or Right Column** list.
2. To position a panel above or below the other panels, click **Move Up or Move Down**.
3. To move a panel to the other side of the user's bookmarks page, click **Move > or < Move**.

NOTE: The system displays all panels under Bookmarks Panel Arrangement for all licensed features regardless of whether or not you enable the corresponding feature for the role.



The maximum number of combined bookmarks a role can have is approximately 500. If a role has more than 500 bookmarks, some operations (for example, delete role, duplicate role) may not work correctly. The workaround is to split a role with a large number of bookmarks into multiple roles.

-
4. Under **Help Page**, select **options** to control the **Help page** that appears when users click the Help button on the toolbar:
 - **Disable help link** - Select this option to prevent users from displaying Help by removing the Help button from the toolbar.
 - **Standard help page** - Select this option to display the standard end-user Help.

- **Custom help page** - Select this option to display a custom Help page. Specify the URL to the custom help page, and then provide an optional width and height for the help page's window. The system rewrites the URL and creates an access control rule to allow users access to the URL. (Note that users can also enter the custom URL in the Browse field on the toolbar.) The system evaluates the access control rule after all other policies, which means another policy could deny access to the URL. (Note that when you choose this option, the system disables the Tips link next to the Browse field.)
 - **Also allow access to directories below this url** - Select this option to allow users access to subdirectories of the custom help page URL. For example, if you specify `http://www.domain.com/help`, users can also access `http://www.domain.com/help/pdf/`.
5. Under User Toolbar, select options for the toolbar on the Bookmarks page and other secure gateway pages:
- **Home** - Select this option to display the Home icon on the Bookmarks page and other secure gateway pages.
 - **Preferences** - Select this option to display the Preferences button.
 - **Session Counter** - Select this option to display a time value on the user toolbar that indicates the maximum remaining time allowed in the user's current session. Note that a period of user inactivity could also end the current session before this maximum time expires.
 - **Client Application Sessions** - Select this option to display the Client Apps button on the user toolbar. Users can click this button to display the Client Application Sessions page where they can start client applications such as VPN Tunneling or Secure Application Manager. If you do not select this option, the system displays the Client Application Sessions panel on the Bookmarks page.
6. Under Browsing toolbar, select options for the toolbar that users see when browsing pages not located on the system, such as external web sites:
- **Show the browsing toolbar** - Select this option to display the browsing toolbar.
 - **Toolbar type** - Select the type of browsing toolbar you want to display:
 - **Standard** - This toolbar can be moved to the top left or top right side of the browser window. Users can also collapse and expand the toolbar. When collapsed, the toolbar displays the custom logo only. The toolbar's default state is expanded and on the top right side of the browser window.
 - **Framed** - This toolbar remains fixed in a framed header section at the top of the page.

We recommend that you do not use the top variable when working with a frame set because after the system intermediates the page, top might reference a different frame than you intend.



This change might make the framed toolbar disappear or could cause your intermediated application to work erratically or incorrectly. See Content Intermediation Engine Developer Guide,

- **Toolbar logo and Toolbar logo (mobile)** - Specify a custom logo (such as your company's logo) that you want to display on the standard and framed toolbars by browsing to the image file (optional). When the user clicks the logo, the page you specify for the Logo links to option appears. The current logo for the browsing toolbar appears next to these options.
- **Logo links to** - Select an option to link the browsing toolbar logo to a page that appears when users click the logo:
 - **Bookmarks page** - Links the logo to the Bookmarks page.
 - **Start Page" settings** - Links the logo to the custom start page you specified under the Start Page section. In the welcome message of the sign in page, the admin can now include hyperlinks with VMWare-View custom protocol (vmware-view://). Therefore the set of allowed hyperlinks are now vmware-view, http, https, mailto, ftp.
 - **Custom URL** - Links the logo to the URL you enter in the associated text box (optional). This resource must be accessible to the system. The system rewrites the URL and creates an access control rule to allow users access to the URL. (Note that users can also enter the custom URL in the Browse field on the toolbar.) The system evaluates the access control rule after all other policies, which means another policy could deny access to the URL.
 - **Also allow access to directories below this url** - Select this option to allow users access to subdirectories of the custom URL.
- Specify the items you want to display in the browsing toolbar:
 - **Enable "Home" link** - Select this option to display the Home Page button, which is linked to the Bookmarks page.
 - **Enable "Add Bookmark" link** - Select this option to display the Bookmark this Page button.

- **Enable "Bookmark Favorites" link** - Select this option to display the Bookmark Favorites button. When the user clicks this button, the system displays a list of the bookmarks that the user specified as favorites on the Add Web Bookmark page of the secure gateway.
- **Display Session Counter** - Select this option to display a time value on the browsing toolbar that indicates the maximum remaining time allowed in the user's current session. Note that a period of user inactivity could also end the current session before this maximum time expires.
- **Enable "Help" link** - Select this option to display the Help button, which is linked to the Help page you specify for under Help page.



If you click **Users > User Roles > Role Name > Web > Options** and clear the User can add bookmarks check box, then the system does not display the Bookmark this Page and Bookmark Favorites buttons on the browsing toolbar even if you select the Enable "Add Bookmark" link and Enable "Bookmark Favorites" link options.

- **Use Iframe in Toolbar** - Select this option if you are having problems with using iframes with JavaScript rewriting and with the Firefox web browser. This option resolves interoperability problems with the above.
7. Under Personalized greeting, specify a greeting and notification message on the Bookmarks page (optional):
- **Enabled** - Select this option to display the personalized greeting. The system displays the username if the full name is not configured.
 - **Show notification message** - Select this option and enter a message in the associated text box (optional). The message appears at the top of the Bookmarks page after you save changes and the user refreshes that page. You may format text and add links using the following HTML tags: `<i>`, ``, `
`, `` and `<a href>`. However, the system does not rewrite links on the sign-in page (because the user has not yet authenticated), so you should only point to external sites. Links to sites behind a firewall will fail. You may also use Ivanti Connect Secure system variables and attributes in this field.



The length of the personalized greeting cannot exceed 12K, or 12288 characters.

If you use unsupported HTML tags in your custom message, the system may display the end user's home page incorrectly.

8. Under **Other**, specify whether or not you want the copyright notice and label shown in the footer (optional). This setting applies only to those users whose license permits disabling the copyright notice. For more information about this feature, call Support.
9. Click **Save Changes**. The changes take effect immediately, but current user browser sessions may need to be refreshed to see the changes.
10. Click **Restore Factory Defaults** to reset all user-interface options back to factory defaults (optional).

Optimized Interface for the Apple iPad

The system is optimized for a number of platforms, including the Apple iPad. This optimization includes:

- **Login pages** - includes the login and logout pages as well as intermediate pages that appear after the user enters their credentials on the sign-in page and before the Home page appears. The following is a list of these customized login pages:
 - Cancel.shtml
 - Defender.shtml
 - ExceededConcurrent.shtml
 - GeneratePin.shtml
 - GraceLoginUsed.shtml
 - LoginPage.shtml
 - Logout.shtml
 - NewPin.shtml
 - NextToken.shtml
 - PasswordChange.shtml
 - PasswordExpiration.shtml
 - SelectRole.shtml
 - ShowSystemPin.shtml
 - SigninNotifPostAuth.shtml
 - SigninNotifPreAuth.shtml

- SM-NewPinSelect.shtml
- SM-NewPinSystem.shtml
- SM-NewUserPin.shtml
- SM-NextToken.shtml
- SSL.shtml
- confirmation.shtml
- confirmation_opensessions.shtml
- user_unknown.shtml
- **Home page** - This home page displays the welcome panel and any applicable notification messages as well as the Web Bookmark panel, the File Bookmark (or Files) panel, the VPN and Preferences button.
- **Web Bookmark pages** - Located on the home page, the Web Bookmark panel lists each individual bookmarks and allows user to tap and browse the bookmark destination page. To edit bookmarks, tap the Edit button on the panel header and the Edit Bookmark page appears. On this page, user can edit individual bookmarks, reorder bookmarks, and delete bookmarks. Editing is limited to user-created bookmarks.
- **File Bookmark pages** - Located on the home page, the File Bookmark panel lists each individual bookmarks. To edit bookmarks, tap the Edit button on the panel header and the Edit File Bookmark page will be displayed. On this page, user can edit individual bookmarks, reorder bookmarks, and delete bookmarks. Editing is limited to user-created file bookmarks.
- **Preferences page** - Located the home page is a Preferences button. When tapped, it displays the Preferences setting page, containing configuration options for changing username, delete cookies, delete session cookies and delete passwords.
- **Error pages** - Error pages that can be seen while using the features made available on the iPad are customized.
- **Company logos** - Most pages display a company logo. These pages are capable of displaying custom logos if uploaded from the admin GUI.

The following table lists the supported configurable options on the Apple iPad:

Custom User Interface Options	Supported
Header Logo Image	Yes
Header Background Color	No
Sub-Header Background Color	No
Sub-Header Text Color	No
Start Page Message (Welcome message)	Yes
Bookmark Panel Arrangement	No
Enable/Disable Help Link	Yes
Window Size of Help Page	No
Show/Hide Preferences Toolbar	No
Show/Hide Session Counter	Yes
Browsing Toolbar Items	Yes
Post-Auth Sign-In Notification	Yes
Personalized Greetings	Yes
Show Copyright Notice in Footer	No

Defining Default Options for User Roles

You can define default options for all user roles, just as you can for delegated administrator roles. Default values are used for newly created roles or for roles where the session or UI option check boxes are not selected in the User > User Roles > *UserName* > General > Overview window.

The default options include, but are not limited to:

- **Session Options**
 - **Session lifetime** - Define the idle timeout, maximum session length, and reminder time in minutes.
 - **Enable session timeout warning** - Determine whether to display warning and login page.
 - **Roaming Session** - Define level of mobility access.
 - **Persistent Session** - Define state across browser instances.

- **Persistent password caching** - Define password state across sessions.
- **Browser request follow-through** - Define response to browser session expiration.
- **Idle timeout application activity** - Define system response to application session activity.
- **UI Options**
 - **Header** - Define the logo and background color.
 - **Sub-headers** - Define the background and text color.
 - **Start page** - Define which page appears after the user logs in.
 - **Bookmarks Panel Arrangement** - Define the panels that appear on the user's bookmark page.
 - **Help Page** - Display standard or custom help.
 - **User Toolbar** - Define the links that appear on a user's home page.
 - **Browsing toolbar** - Define the links that appear when a user is browsing an external web site.
 - **Personalized Greeting** - Display user's name and notification message on the user's welcome page.
 - **Bookmarks Panel Arrangement Other** - Show copyright notice.

Defining Default Options for User Roles

To define the default options for all user roles:

1. Select **Users > User Roles**.
2. Click **Default Options**.
3. Modify settings in the **Session Options, UI Options, and Custom Messages** tabs.
4. Click **Save Changes**. These become the new defaults for all new user roles.

If you do not want user roles to see the copyright notice, you can also clear the Show copyright notice and "Secured by" label in footers check box for user roles, in general. That way, all subsequent roles you create do not allow the notice to appear on the end user UI.

Customizing Messages

You can customize basic messages that may be displayed to your end users when they sign in to the device. You can change the message text, and you can add internationalized versions of the messages in Chinese (Simplified), Chinese (Traditional), French, German, Japanese, Korean, and Spanish, in addition to English.

To customize messages:

1. Select **Users > User Roles**.
2. Click **Default Options**.
3. Select the **Custom Messages** tab.
4. Select the language to use from the menu.
5. Enter your text in the **Custom Message** box, below the default message you want to override.
6. Click **Save Changes**.
7. Repeat the process to create messages in additional languages.

Customizing UI Views for User Roles

You can use customization options on the Roles page to quickly view the settings that are associated with a specific role or set of roles. For instance, you can view all of the user roles and any Web bookmarks that you have associated with them. Additionally, you can use these customized views to easily link to the bookmarks and other configuration settings associated with a role.

To view a sub-set of data on the Roles page:

1. Click **Users > User Roles**.
2. Select an option from the View list at the top of the page. Table 5 describes these options.
3. Select one of the following options from the For list:
 - **All roles** - Displays the selected bookmarks for all user roles.
 - **Selected roles** - Displays the selected bookmarks for the user roles you choose. If you select this option, select one or more of the check boxes in the Role list.
4. Click **Update**.

The following table lists the View Menu Options

Option	Description
Enabled Settings	Displays a graph outlining the remote access mechanisms and general options that you have enabled for the specified roles. Also displays links (the check marks) that you can use to access the corresponding remote access and general option configuration pages.
Restrictions	Displays Host Checker restrictions that you have enabled for the specified roles. Also displays links you can use to access the corresponding Host Checker configuration pages.
VPN Tunneling	Displays VPN Tunneling settings that you have configured for the specified roles. Also displays links you can use to access the corresponding VPN Tunneling configuration pages.
Role Mapping Rule & Realms	Displays the assigned authentication realms, role mapping rule conditions, and permissive merge settings for the specified roles. Also displays links you can use to access the corresponding realm and role mapping configuration pages.
Bookmarks: All	Displays the names and types of all of the bookmarks that you have enabled for the specified roles. Also displays links you can use to access the corresponding bookmark configuration pages. (Note that if you created a bookmark through a resource profile, the link appears in the Resource column. Otherwise, the link appears in the Bookmark column.)
Bookmarks: Web	Displays the Web bookmarks that you have enabled for the specified roles. Also displays links you can use to access the corresponding bookmark configuration pages. (Note that if you created a bookmark through a resource profile, the link appears in the Resource column. Otherwise, the link appears in the Web Bookmark column.)
Bookmarks: Files (Windows)	Displays the Windows File bookmarks that you have enabled for the specified roles. Also displays links you can use to access the corresponding bookmark configuration pages. (Note that if you created a bookmark through a resource profile, the link appears in the Resource column. Otherwise, the link appears in the Windows File Bookmark column.)
Bookmarks: Terminal Services	Displays the Terminal Services bookmarks that you have enabled for the specified roles. Also displays links you can use to access the corresponding bookmark configuration pages. (Note that if you created a bookmark through a resource profile, the link appears in the Resource column. Otherwise, the link appears in the Terminal Services Session column.)

Option	Description
ACL Resource Policies: All	Displays the resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.
ACL Resource Policies: Web	Displays the Web resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.
ACL Resource Policies: Files (Windows)	Displays the Windows file resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.
ACL Resource Policies: SAM	Displays the JSAM and PSAM resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.
ACL Resource Policies: Terminal Services	Displays the Terminal Services resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.
ACL Resource Policies: VPN Tunneling	Displays the VPN Tunneling resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.
Resource Profiles: All	Displays the resource profiles that are associated with the specified roles. Includes the type, name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.
Resource Profiles: Web Applications	Displays the Web application resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.

Option	Description
Resource Profiles: Web Hosted Java Applets	Displays the hosted Java applet resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.
Resource Profiles: Files (Windows)	Displays the Windows file resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.
Resource Profiles: SAM Client Applications	Displays the JSAM and PSAM application resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.
Resource Profiles: SAM PSAM destinations	Displays the PSAM destination resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.
Resource Profiles: Terminal Services	Displays the Terminal Services resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.

Virtual Desktop Resource Profiles

Virtual Desktop Resource Profile Overview

In addition to standard resource profiles and resource profile templates, you can configure virtual desktops as resource profiles.

As with the other resource profiles, a virtual desktop profile contains all of the role assignments and end-user bookmarks required to provide access to an individual resource. Unlike other resource profile types, there is no resource policy to configure for virtual desktops due to the dynamic nature of virtual desktops. The IP address and port of the system is not known until the end user launches a session so dynamic ACLs are used.

Icons in the Virtual Desktops section on the end user's home page represent desktops defined by the administrator. Clicking the icon launches the session using the Virtual Desktop Infrastructure (VDI) architecture.

A few of the main features of virtual desktop resource profiles are:

- SSO so that the user can sign on without having to enter their credentials
- Dynamic ACLs
- Client delivery mechanism for end users who do not have the client already installed on their system
- Connection logging

Configuring a Citrix XenDesktop Resource Policy

The Citrix XenDesktop manages a pool of virtual desktops hosted on virtual machines and provides the connection management to those desktops. A list of XenDesktops is displayed to the end user as bookmarks. When a desktop is selected, the Citrix client is launched and the user can access that desktop.

To configure a Citrix XenDesktop profile:

1. Select **Users > Resource Profiles > Virtual Desktops**.
2. Click **New Profile**.
3. Select **Citrix XenDesktop** from the Type drop-down list.
4. Enter a name and description (optional) to identify this profile.

5. Enter the name or IP address and port of the connection broker using the format ip:port. For example,

10.10.1.10:80

xml.example.com:80

You can enter more than one IP address. Place each address on a separate line.

6. Select the **Use SSL for connecting to the Server** check box if **SSL** is required to connect to the server.
7. Enter the username to connect to the connection broker or use the **<USERNAME>** session variable.
8. Enter the password:
 - To use a variable password to connect to the connection broker, select **Variable Password** and enter the variable in the form of **<PASSWORD>** or **<PASSWORD@SEcAuthServer>**.
 - Select **Password** to use a static password to connect to the connection broker and enter the user credential's password.
9. Enter the domain where the connection broker is located.
10. Select **Enable Java support** to specify a Java applet to use to associate with the resource profile. The system uses this applet to intermediate traffic or falls back to this applet when ActiveX is not available on the user's system.
11. Click **Save** and **Continue**.
12. Select the roles to which this profile applies and click **Add**.
13. The Enabled Settings table under Users > User Roles also displays which roles have virtual desktops enabled.
14. Click **Save Changes**.
15. (**Optional**.) In the Bookmarks tab, modify the default bookmark created by the system and/or create new ones.

Configuring a VMware View Manager Resource Profile

VMware View Manager, formerly VMware VDI, lets you run virtual desktops in a data center that provide end users a single view of all their applications and data in a personalized environment regardless of the device or location they log in from.

To configure a VMware View Manager profile:

1. Select **Users > Resource Profiles > Virtual Desktops**.
2. Click **New Profile**.
3. Select **VMware View Manager** from the Type drop-down list.
4. Enter a name and description (**optional**) to identify this profile.
5. Enter the **name** or **IP address** and port of the connection broker using the format ip:port. For example,

10.10.1.10:80

xml.example.com:80

You can enter more than one IP address. Place each address on a separate line.

6. Select the **Use SSL for connecting to the Server** check box if **SSL** is required to connect to the server.
7. Enter the username to connect to the connection broker or use the **<USERNAME>** session variable.
8. Enter the password:
 - To use a variable password to connect to the connection broker, select **Variable Password** and enter the variable in the form of **<PASSWORD>** or **<PASSWORD@SEcAuthServer>**.
 - Select **Password** to use a static password to connect to the connection broker and enter the user credential's password.
9. Enter the domain where the View Manager server is located.
10. Click **Save** and **Continue**.
11. Select the roles to which this profile applies and click **Add**.

12. The Enabled Settings table under Users > User Roles also displays which roles have virtual desktops enabled.
13. Click **Save Changes**.
14. (Optional.) In the Bookmarks tab, modify the default bookmark created by the system and/or create new ones.

Defining Bookmarks for a Virtual Desktop Profile

When you create a virtual desktop resource profile, the system automatically creates a bookmark that links to the server that you specified in the resource profile. The system allows you to modify this bookmark as well as create additional bookmarks to the same server.

These bookmarks are listed in the role bookmark pages (Users > User Roles > Role_Name > Virtual Desktop > Sessions) but you cannot add, modify or delete the bookmarks from the role bookmarks page. Bookmarks can only be added as part of the resource file.

To configure resource profile bookmarks for virtual desktop profiles:

1. Select **Users > Resource Profiles > Virtual Desktop**.
2. Click the name of the virtual desktop profile.
3. Click the Bookmark tab to modify an existing session bookmark. Or, click **New Bookmark** to create an additional session bookmark.
4. (Optional.) Change the name and description of the session bookmark. (By default, the system populates and names the session bookmark using the resource profile name.)
5. Specify whether all desktops or to a selected subset of desktops are available to the user.

The desktop list is retrieved from the connection broker using the credentials defined in the profile resource page.

6. Enter the credentials used to log in to the actual VMware or XenDesktop machine. The system passes these credentials to the server so that users can sign on without having to manually enter their credentials.

7. Specify how the window should appear to the user during a session by configuring options in the Settings area of the bookmark configuration page.

(XenDesktop) Under Preferred Client, you can select Automatic Detection, Citrix Client or Java. If you select Automatic Detection, the system checks to see if Citrix Client is present. If it is not present, the end user is given the choice to download the Citrix Client or to use the alternate client, Java ICA Client.

8. Allow users to access local resources such as printers and drives through the terminal session by configuring options in the Connect Devices area of the bookmark configuration page.

(VMware) Enable **MMR** - Redirect certain multimedia codecs running on the remote desktop to the local client for rendering of full-motion video and audio.

(VMware) **Allow Desktop Reset** - Allow users to reset their desktop without administrative assistance. For example, if the desktop hangs, there is currently no way for the user to perform a hard reboot of the desktop. This option allows the users to restart their own virtual desktops thereby reducing the dependency on the administrator or helpdesk.

9. Specify how the terminal emulation window should appear to the user during a terminal session by configuring options in the Desktop Settings area.
10. Specify the roles to which you want to display the session bookmarks if you are configuring the session bookmark through the resource profile pages, under Roles:
 - **ALL selected roles** - Displays the session bookmark to all of the roles associated with the resource profile.
 - **Subset of selected roles** - Displays the session bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL Selected Roles list and click Add to move them to the Subset of selected roles list.
11. Click **Save Changes**.

Configuring the Client Delivery

You can use the Virtual Desktop Configuration page to define the client delivery mechanism for end-users who do not have the client. The process is similar for both Citrix XenDesktop and VMware View Manager.

1. Choose **System > Configuration > Virtual Desktops**. For Citrix XenDesktop, select **Citrix**.

2. Select **Download from Ivanti Connect Secure** to download the client file from the system. Click Browse to locate the client file (.msi, .exe or .cab) and enter the version number.
3. Select **Download from a URL** to download the client file from the Internet. If desired, enter a new URL to override the default.
4. Check the **Access the URL through the Ivanti Connect Secure** check box if end users cannot directly access the specified Web page. Selecting this option allows users to use the secure gateway to access the URL.
5. Under Server Connection Timeout, enter the number of seconds to wait for the server to respond before timing out.

Connecting to the Servers

When an end user clicks a desktop icon, the system passes credentials to the server based on the desktop profile.

For XenDesktop, the system authenticates to the Citrix DDC server using credentials defined in the desktop profile. If successful, the list of available desktops is returned by the DDC server and is represented as bookmarks to the end user. When an end user clicks a XenDesktop icon, the system retrieves the ICA from the XenDesktop server and presents a desktop session to the user.

When an end user clicks a VMware View Manager icon, the system authenticates to the View Manager using credentials defined in the desktop profile. If authentication is successful, a JSESSIONID cookie is returned by the View Manager, the system creates a tunnel using the cookie for the duration of the session.

If the desktop is unavailable, the client will continue to try to connect until the desktop is available or until a predefined timeout period occurs. An error message lets the user know the status, either that the system is retrying the connection or that the desktop is unavailable. Similarly if the desktop is already in use by another enduser, an error message is presented to the user.

User logs are updated to show which VM machines are assigned to each user. Username, realm, VM IP, port, connection type, pool and connection broker are logged with each message.

The Active Virtual Desktops Sessions page (System > Status > Virtual Desktop Sessions) lists the active connections, including the connection broker, the VM machine assigned to the user and the connection type.

Authentication and Directory Servers

AAA Server Overview

Understanding the Role of AAA Servers in the Ivanti access management framework



From 9.1R15 onwards, support for Siteminder, NIS, Anonymous, Novell eDirectory, and iPlanet AAA servers are deprecated. Ensure you remove all configurations related to these servers before upgrading to 9.1R15. **Upgrade may fail if all configurations are not removed.** For more information refer [KB45044](#).



When importing the configurations from 9.1R14 and prior releases, authentication server configurations are imported as is, but the user logins may not work. The deprecated authentication servers are highlighted in Red.

AAA stands for authentication, authorization, and accounting. A AAA server is a database that stores user credentials - username and password - and, in some cases, group information or other user attributes. The authentication results and the group or user attribute information is used by the Ivanti access management framework for policy decisions.

In the Ivanti access management framework, the sign-in page, realm, and AAA server configurations are associated. They determine user access and user role. A user submits credentials through a sign-in page, which specifies a realm, which is associated with a AAA server. If the access request meets the realm's authentication policy, the system forwards the user's credentials to the associated authentication server. The authentication server's job is to verify the user's identity. After verifying the user, the authentication server sends approval. If the realm also uses the server as a directory/attribute server, the AAA server sends the user's group information or other user attribute information. The access management framework then evaluates the realm's role-mapping rules to determine the user roles that apply to the session.

The Ivanti access management framework supports the following types of AAA servers:

- Local - You can create special purpose local databases to manually create user accounts, or manage access based on digital certificates.
- External (standards-based) - You can integrate standards-based LDAP and RADIUS servers with the access management framework. In addition to using the backend server for authentication, you can use LDAP group and RADIUS attribute information in role-mapping rules.

- External (other) - You can integrate compatible versions of popular third-party AAA servers with the access management framework. In addition to using the backend server for authentication, you can use Active Directory group information in role-mapping rules. In addition, you can use MDM device attributes in role mapping rules.

The following table is a reference of the AAA servers supported in Ivanti Connect Secure deployments.

The following table lists the Supported AAA Servers:

Ivanti Connect Secure	
Local	"Local Authentication Server"**, "Certificate Server", "SAML Server"*** **No special features to manage guest users. ***Supports an authentication server configuration when deployed as a SAML service provider. Different Ivanti Connect Secure features support a local SAML server when deployed as a SAML identity provider.
External (standards-based)	"LDAP Server", "RADIUS Server"
External (other)	"Active Directory", "MDM Server", "RSA ACE Server"

AAA Server Configuration Task Summary

To integrate an authentication server:

1. Configure the authentication server. Select **Authentication > Authentication > Auth. Servers page** and complete the authentication server configuration.
2. Create an authentication realm. Select **Users > User Realms or Administrators > Admin Realms** and select the authentication server when you complete the authentication realm configuration.

AAA Traffic Management

From 9.0R3 release, the Ivanti Connect Secure Virtual appliances and the Ivanti Secure Appliances allow the administrator to choose the communicating interface or the network for each authentication server.

This feature allows the AAA traffic across the following interfaces:

- Physical Internal
- Physical External

- Physical Management
- Virtual ports for Physical Interfaces
- VLAN ports
- Virtual Ports on VLAN Interfaces


This feature allows to connect to remote supported authentication servers through any interfaces based on the network Topology.

The following Authentication server types are supported:

- LDAP
- Active Directory
- RADIUS
- CRL and OCSP traffic flow


Configuring AAA Traffic Management Across Interfaces

1. Select **Authentication > Auth Servers** and configure service provider AAA configurations as needed.



The screenshot shows the 'Authentication Servers' configuration page. At the top, there is a green title 'Authentication Servers'. Below the title is a blue button labeled 'Enable Auth Traffic Control'. Underneath, there is a 'New:' label followed by a dropdown menu with '(Select server type)' and a downward arrow. To the right of the dropdown are two blue buttons: 'New Server...' and 'Delete...'. At the bottom left, there is a dropdown menu showing '10' and a downward arrow, followed by the text 'records per page'.

2. Click **Enable Auth Traffic Control**. A new window appears.



The screenshot shows a dialog box titled 'Confirm Traffic Decoupling Enabling'. At the top left, there is a green triangle icon followed by the title. Below the title is a yellow warning box with a triangle icon and the text: 'Warning: Enabling Traffic Decoupling can lead to user authentication failure if Auth Servers are not reachable via configured interfaces'. At the bottom, there are two blue buttons: 'Enable Traffic Decoupling' and 'Cancel'.

- Click **Enable Traffic Decoupling** to confirm. The page navigates to the Auth server page that displays the options to configure the AAA traffic interfaces.

Authentication Servers

▼ **NIC Selection**

Global Setting
 Auth Server Level

Chose the interface from the below list for reaching to Auth Server

localhost2 Management management 2.2.144.33

Save Disable Auth Traffic Control

- Select **Global setting** to use same interface across all supported authentication servers or select **Auth Server Level** to select the interface for a specific authentication server for the AAA traffic.

Authentication Servers

▼ **NIC Selection**

Global Setting
 Auth Server Level

Chose the interface from the below list for reaching to Auth Server

PCS_34 External external 2.2.144.34

Save Disable Auth Traffic Control

New: (Select server type) New Server... Delete...

10 records per page Search:

Authentication/Authorization Servers	Type
Administrators	Local Authentication
AD-Pbuindia.com	Active Directory / Windows NT
Cert-Auth	Certificate Server
LDAP-Server	LDAP Server
NIS_Server	NIS Server
RadiusSBR	RADIUS Server
Steminder	SiteMinder Server
System Local	Local Authentication

← Previous 1 Next →

- Select the required interface and port from the list.
For Clusters, select applicable interfaces and associated ports.
- Click **Save**.


Upgrading from Previous Releases

Prior to 9.0R3, the AAA traffic was routed through the management port. This option was available on both the Default Network and the Administrative Network.

On upgrading to 9.0R3, the AAA traffic can be routed through Internal, External, Management, Virtual ports and VLAN ports. If Send AAA traffic via Management Port was enabled in pre-9.0R3, then by default, immediately after upgrade, the AAA traffic is routed through the management port for all authentication servers as a global setting. The selected interfaces may be modified as required using the Global Settings or the Auth Server Level settings.

Configuring AAA Traffic Management on Upgrade

Prior to 9.0R3 release, the Ivanti Connect Secure Interface had the option "*Send AAA traffic via Management Port*" as shown in figure under System > Traffic Segregation. This option routes the AAA traffic through the Management port by default.



On upgrade to 9.0R3 release, if Send AAA traffic via Management Port was enabled in pre-9.0R3, then the following changes are observed:

- The AAA traffic management options are available under Authentication > Auth. Servers.
- The AAA traffic management is enabled by default.
- The physical port is automatically set to Management port or default VLAN.

Authentication Servers

▼ NIC Selection

Global Setting Auth Server Level

Chose the interface from the below list for reaching to Auth Server

localhost2 Management management 2.2.144.33

Save Disable Auth Traffic Control

Using the Local Authentication Server

This topic describes the local authentication server.

Local Authentication Server Overview

This section provides an overview of the feature and its limitations.

Understanding the Local Authentication Server

The local authentication server is an authentication database that is built in to the system. Therefore, it is considered a "local" server in contrast to a third-party enterprise AAA server that is connected over the network.

Typically, you create local user accounts for temporary users who do not have accounts on your enterprise AAA servers. Temporary users include lab users or guests, but you might find the local authentication server useful to create temporary accounts for users who are normally verified by an enterprise AAA server that you plan to disable.

You also use the local authentication server to create accounts for administrator users, such as system administrators.



Although it is common practice to use the local authentication server for administrator accounts, it does not preclude you from using any of the supported third-party enterprise AAA servers in your administrator access management framework.

Configuring the Local Authentication Server

You can create multiple local authentication server instances. When you define a new local authentication server, you must give the server a unique name and configure options for passwords.

To create a local authentication server:

1. Select **Authentication > Auth. Servers**.
2. Select Local Authentication and click **New Server** to display the configuration page.
3. Complete the configuration as described in the following table.
4. Save the configuration.

The following table lists the Local Authentication Server Configuration Guidelines:

Settings	Guidelines
Name	Specify a name that is useful to you.
Password Options	
Minimum length	Specify a number of characters. The valid range is 0-99. 6 is the default.
Maximum length	Specify a number of characters. The valid range is 0-99. 8 is the default. The maximum length cannot be less than the minimum length.
Minimum digits	Specify the number of digits required in a password. Do not require more digits than the value of the maximum length option.
Minimum letters	Specify the number of letters required in a password. Do not require more letters than the value of the maximum length option. If you enable the previous option, the combined total of the two options cannot exceed that of the value specified in the maximum length option.
Uppercase and lowercase required	Select this option if you want all passwords to contain a mixture of uppercase and lowercase letters. Require passwords to contain at least two letters if you also require a mix of uppercase and lowercase letters.
Different from username	Select this option if the password cannot equal the username.
Different from previous password	Select this option if a new password cannot equal the previous password.
Stored as cleartext	Select this option if you are using open authentication protocol sets. CHAP and EAP-MD5-Challenge work with local authentication servers only if you select this option. Be aware of the security implications of storing passwords as cleartext.
Password Management	

Settings	Guidelines
Allow users to change passwords	Select this option if you want users to be able to change their passwords. In addition to selecting local authentication password management options, you must select the Enable Password Management option for the associated realm authentication policy.
Force password change	Select this option to specify the number of days after which a password expires. The default is 64 days.
Prompt users to change password	Select this option to specify when to prompt the user to change passwords.
Account Lockout	
Enable account lockout for users	Select this option to manage user authentication failures for admin users of local authentication server.
Maximum wrong password attempts	Specify the number of consecutive wrong password attempts after which the admin user account will be locked. The default value is 3 retries.
Account Lockout period	Specify the time in minutes for which admin user account will remain locked. The default value is 10 minutes.

Creating User Accounts

You use the Users page to create local authentication server user accounts. A user account includes a username and password to be used for authentication, as well as other information used for records and account management.

To create a local user account:

1. Select **Authentication > Auth. Servers**.
2. Select the local authentication server to which you want to add a user account.
3. Click the **Users** tab.
4. Click **New** to display the configuration page.
5. Complete the configuration as described in the following table.
6. Save the configuration.

The following table lists User Account Configuration Guidelines:

Settings	Guidelines
Username	Do not include "~" in a username. You cannot change a username after you create the account.
Full Name	Specify the user's full name.
Password	Specify a password. Make sure that the password you enter conforms to the password options specified on the local authentication server configuration page.
Confirm password	Confirm the password.
One-time use	Select this option to limit the user to one login. After one successful login, the user's login state is set to disabled, and the user receives an error message when attempting subsequent sign-ins. However, you can manually reset this option to allow the same user to log in again.
Enabled	Select this check box if it is not already selected. If the one-time use option has been implemented, this option is listed as Disabled after the user has logged in successfully. If a permanent or one-time user is logged in and you disable this option, the user is immediately logged out of the system and receives an error message.
Require user to change password	Select this option to force users to change their passwords at the next login. If you force the user to change passwords, you must also enable the local authentication password management options.

Managing User Accounts

You use the Users page to list, modify, and delete local authentication server user accounts.

The screenshot shows a web interface for managing users. At the top, there are tabs for 'Settings', 'Users' (which is active), and 'Admin Users'. Below the tabs, there is a search field labeled 'Show users named:' with an asterisk (*) entered, a 'Show' dropdown set to '200 users', and an 'Update' button. Below this are three buttons: 'New...', 'Delete...', and 'Unlock...'. A pagination control shows 'Page 1 of 1' with navigation arrows. The main content is a table with the following structure:

	Username	Name	Last Sign-in Statistic			
			Date&Time	IPAddress	Agent	Status
<input type="checkbox"/>	lu1	Unspecified Name	2019/03/04 12:23:46	172.21.24.57	Mozilla/5.0 (Macintosh; Intel Mac OS X 10_13_3) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/72.0.3626.119 Safari/537.36	
<input type="checkbox"/>	lu2	Local User 2	2019/03/05 03:43:35			
<input type="checkbox"/>	lu3	lu3	2019/03/05 03:43:49			

To manage a user account:

1. Select **Authentication > Auth. Servers**.
2. Click the link for the authentication server you want to manage.
3. Click the **Users** tab to display the user accounts table.

The user accounts table includes entries for the accounts that have been created. The Last Sign-in Statistic column shows the last successful sign-in date and time for each user, the user's IP address, and the agent or browser type and version. The Status column for the user shows the account-locked warning icon if the user account is locked.

4. Use the controls to search for users and manage user accounts:
 - To search for a specific user, enter a username in the Show users named box and click **Update**.

You can use an asterisk () as a wildcard, where * represents any number of zero or more characters. For example, to search for all usernames that contain the letters jo, enter *jo*. The search is case-sensitive. To display the entire list of accounts again, type * or delete the field's contents and click Update.*

- To limit the number of users displayed on the page, enter a number in the Show N users box and click Update.
- To edit the user account configuration, click the link in the Username column to display the Update Local User Account page.

- To terminate the user session and delete the account, select the box next to the user account record and click Delete.
- To unlock a user account, select the locked-out account and click Unlock. The account-locked warning icon will disappear after successful unlock.
- To view the admin user access logs, select **System > Log/Monitoring > Admin Access > Log**.

Select a user to display the user account configuration page. You can use this page to modify the account settings, or to disable or quarantine the account.

Creating Administrator User Accounts

You use the Admin Users page to create a special admin user account that enables the account holder to manage the local authentication server users table. These special admin users can sign in to a special page that enables them to create, modify, and delete user accounts.


To create a special admin user account:

1. Select **Authentication > Auth. Servers > System Local**.
2. Click the **Admin Users** tab to display the configuration page.
3. Specify a username, select an authentication realm, and click **Add** to create the administrator user.
4. Save the configuration.

Using the Admin User Sign-In Page to Manage the Local Authentication Users Table

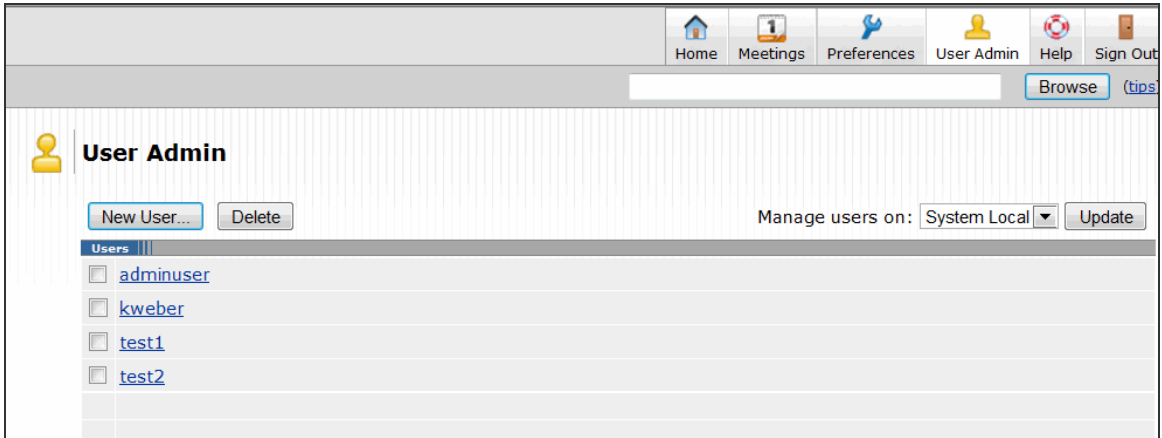
The special admin users created using the feature shown in the previous section can manage the local authentication server accounts table. For example, if an admin user named adminuser is provisioned to manage user accounts for the Users realm, when adminuser signs into the Users realm sign-in page, a User Admin button appears on the toolbar at the top of the page. The following figure depicts shows the toolbar.

The following figure depicts the Sign-In Page Toolbar:



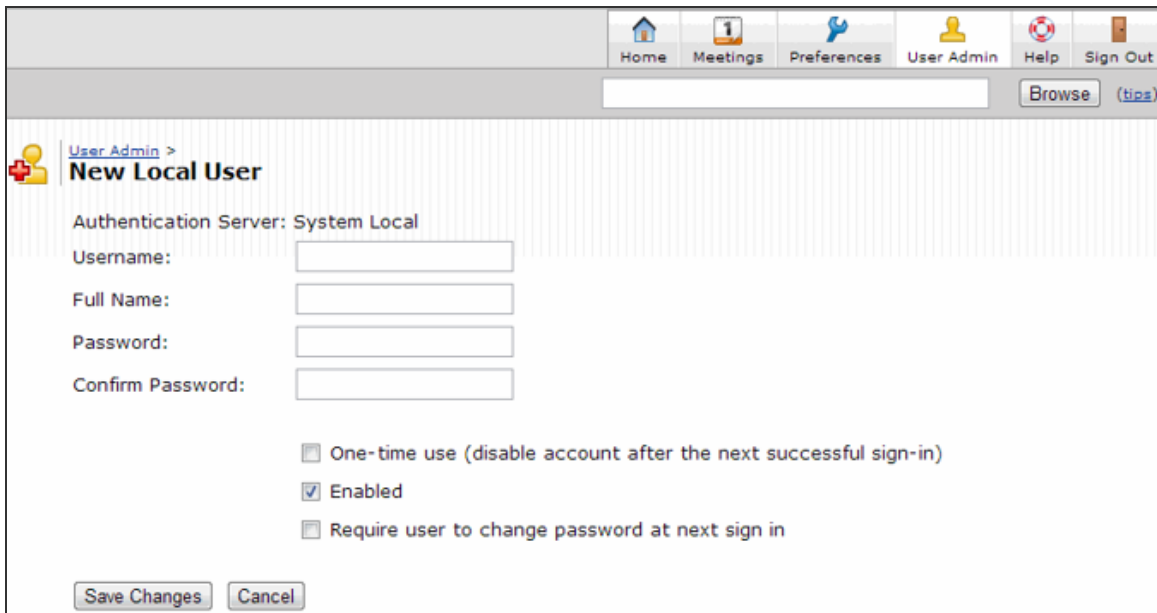
The special admin user can click the User Admin button to display the User Admin page, which shows the local authentication server users table.

The following figure depicts the User Admin Page:



The special admin user can select accounts and delete them and can create user accounts. The same account management guidelines apply as when using the admin console for creating and modifying user records.

The following figure depicts the New Local User Configuration Page:



Using Active Directory

This topic describes integration with the Microsoft® Windows® platform Active Directory™ service.

Microsoft Windows Platform Active Directory Service Overview

This section describes support for using Ivanti Connect Secure with the Active Directory AAA service.

Understanding Active Directory

Active Directory is a directory service used in Windows domain networks. It is included in most Windows server operating systems. Enterprise servers that run Active Directory are called domain controllers. An Active Directory domain controller authenticates and authorizes users and computers in a Windows domain network.

When you use Active Directory as the authentication and authorization service for your Ivanti access management framework, users can sign in to Ivanti Connect Secure using the same username and password they use to access their Windows desktops. You can also use Active Directory group information in role mapping rules.

From 9.1R1 onwards, Active Directory Legacy Mode configuration will not be supported. If you have an existing Active Directory authentication server using Legacy Mode, first migrate to Standard Mode and then upgrade ICS. For the detailed migration procedure, refer KB40430.

Active Directory Feature Support

Ivanti access management framework supports the following Active Directory features:

- Honors trust relationships in Active Directory and Windows NT environments.
- Supports Domain Local Groups, Domain Global Groups, and Universal Groups defined in the Active Directory forest.
- Supports use of Kerberos, NTLMv2, and NTLMv1 authentication protocols.
- Supports user principal name (UPN) format for usernames. This support is available for Web login.

Interoperability Requirements and Limitations

The following limitations apply to interoperability with Active Directory:

- The Ivanti access management framework uses Active Directory security groups, not distribution groups. Security groups allow you to use one type of group for not only assigning rights and permissions, but also as a distribution list for e-mail.

- Each Active Directory configuration you create for the Ivanti access management framework should use a different and unique machine account name.
- If the current Active Directory domain controller is not reachable, the user or machine authentication requests fail for a few seconds (less than 2 minutes) before attempting to authenticate users with another domain controller in the Active Directory domain.
- We do not support Active Directory implementations that use the equal sign operator (=) in a group name, such as: "\=THIRD FLOOR GROUP". The Ivanti access management framework authentication process involves search operations that use the equal sign operator (=) when parsing server catalogs to retrieve group names, usernames and domain names, as well as user_SID and domain_SID values. You might encounter unexpected behavior that can affect normal processing of authentication services if a group name configured on your Active Directory server includes an equal sign operator (=).
- Active Directory versions Windows 2008 R2 and later use a dynamic port range. The default start port is 49152 and the default end port is 65535. Therefore, if there is a firewall between the client service and the Active Directory Service, you must increase the remote procedure call (RPC) port range on the firewall. See *Microsoft Knowledge Base article 929851*.
- The Ivanti password management feature, which enables users to change their Active Directory passwords through the Ivanti service Web server, is not supported for users of trusted domains that do not trust the domain specified in the Ivanti Active Directory configuration.

Configuring Authentication and Authorization with Active Directory Service

To configure integration with Active Directory Service:

1. Select **Authentication > Auth. Servers**.
2. Select **Active Directory / Windows NT** and click **New Server** to display the configuration page.
3. Select **Active Directory mode** and complete the configuration as described in [Table](#).
4. Save the configuration.

The following table lists Active Directory Mode Settings:

Settings	Guidelines
Mode	

Settings	Guidelines
	Select Active Directory mode. This table describes Active Directory mode.
Base Configuration	
Name	Specify a name to identify the server within the system.
Domain	Specify the NetBIOS domain name for the Active Directory domain. The system uses DNS to discover domain controllers in the Active Directory forest. It sends authentication requests to the domain controller at the closest site. Ensure that your DNS servers are configured to resolve the Active Directory domain controller fully qualified domain name (FQDN) and service (SRV) records.
Kerberos Realm	Specify the FQDN of the Active Directory domain. For example, if "pulsesecure" is the domain name (NetBIOS name), then pulsesecure.net is the Kerberos realm name.
Domain Join Configuration	
Username	Specify a username that has permission to join computers to the Active Directory domain. Use the "Delegate Control" workflow in Active Directory to assign the following user account permissions to the username or to a group to which the user belongs: Write Write All Properties Change Password Reset Password Validate Write to DNS hostname Read and write DNS host attributes Delete Computer Objects Create Computer Objects
Password	Specify the password for the special user.
Save Credentials	If this setting is not enabled, the credentials entered will be destroyed after successfully joining the domain. This option is useful when managing clusters. For example, you might want to save the credentials for a cluster node you have yet to add. If you do not enable this option, you must manually enter the credentials when you add the new cluster node.

Settings	Guidelines
Container Name	<p>Specify the container path in Active Directory in which to create the machine account. Changing this field triggers a domain rejoin action. The default is Computers, which is a standard container created during installation of the AD server. The AD Computers container is the default location for new computer accounts created in the domain.</p> <p>If desired, you may specify a different container or OU. To specify nested containers, use a forward slash (/) as the container separator. For example: outer OU/inner OU.</p> <p>Do not use backslashes in the path. Using backslashes causes an Invalid DN Syntax error.</p>
Computer Name	<p>Specify the machine account name. The default computer name is derived from the license hardware in the following format: 0161MT2L00K2C0. We recommend the Computer Name string contain no more than 14 characters to avoid potential issues with the AD/NT server. Do not include the '\$' character.</p>
Update Join Status / Reset Join	<p>The following colors are used to indicate status:</p> <p>Gray. The Domain Join action has not been attempted. This is the default status that appears when you are using the page to create a new Active Directory configuration.</p> <p>Yellow. Attempting to join the Active Directory domain. This is the default status that appears after saving configuration settings or when any domain join settings are changed in an existing configuration.</p> <p>Green. The attempt was successful. This status indicates that this server can now be used to authenticate users.</p> <p>Red. The attempt to join the Active Directory domain was not successful. Click Update Join to get the latest join status of nodes. If the status appears persistently red, click Reset Join to reinitiate the domain join process. The Reset Join action requires Active Directory administrator credentials.</p> <p>For cluster nodes, you might need to click Update Join multiple times to obtain the latest join status of nodes.</p> <p>Transient network issues might also cause the join status indicator to appear red. Before restarting the join process, ensure that it is not caused by network issues. Make sure your DNS servers can resolve queries to the Active Directory domain controller and that the Active Directory credentials are valid and have the appropriate permissions.</p>

Settings	Guidelines
Additional Options	
Authentication Protocol	<p>The system attempts authentication using the protocols you have enabled in the order shown on the configuration page. For example, if you have selected the check boxes for Kerberos and NTLMv2, the system sends the credentials to Kerberos. If Kerberos succeeds, the system does not send the credentials to NTLMv2. If Kerberos is not supported or fails, the system uses NTLMv2 as the next protocol in order.</p> <p>Kerberos. Select this option to enable the Kerberos authentication protocol. Kerberos is the most secure method and is required for Kerberos single sign-on authentication. Kerberos must be enabled if you plan to use Pulse client single sign-on or browser-based agentless single sign-on (SPNEGO).</p> <p>Enable NTLM protocol. Select this option to enable NTLM if you plan to use any of the following features:</p> <ul style="list-style-type: none"> Machine authentication using, Pulse client, or Windows native 802.1x supplicants. MS-CHAP-based authentication protocols for any 802.1x supplicants. User password management. Role mapping rules based on group membership.
Trusted domain lookup	<p>Contact trusted domains. Select this option to contact domain controllers of trusted domains directly without proxying authentication requests and group membership checks through the domain controller.</p> <p>If this option is not selected:</p> <ul style="list-style-type: none"> Network contact with trusted domains is not permitted, but pass-through authentication using the primary domain is still permitted. Trusted domain user's group lookup for Kerberos SSO. Trusted domain user's password-based authentication does not work. Only groups from the domain in which this system is a member are available for use in role mapping when a group search is performed in the server catalog window. If you want to restrict trusted domain users and computers from logging in when this option is not selected, you can define a custom expression based on the ntdomain variable and use it in role mapping rules. For example, if Ivanti Connect Secure belongs to the domain named Corporate, you can define a custom expression as ntdomain=Corporate and use the custom expression in the role mapping rule of the realm.

Settings	Guidelines
Domain Connections	<p>Maximum simultaneous connections per domain. Enter the maximum number of simultaneous domain connections (1 to 10).</p> <p>This field specifies the maximum number of simultaneous connections that the auth daemon should open to the domain controller of one domain. A value of greater than 1 can improve the scalability with simultaneous authentication requests. However, this field value should be judiciously used, especially if trusted domain setting is enabled. This value dictates how many authentication processes are created per domain. For example: if the maximum domain connection is configured as 4 and there are 5 trusted domains, there could be as many as $5 \times 4 + 1 = 21$ auth processes. Hence if there are many trusted domains, the domain connection value needs to be controlled by the administrator, failing which there could be too many auth processes created only for AD authentication purpose. By default, this field value is set to 2 if trusted domain setting enabled. If trusted domain is not enabled, then the default value is set to 5. If Contact trusted domains is enabled, a value above 6 may degrade overall system performance.</p>
Machine account password change	<p>Enable periodic password change of machine account. Select this option to change the domain machine account password for this configuration. Change machine password frequency. Specify a frequency in days. For example, every 30 days.</p>
<p>User Record Synchronization This feature is available only on Ivanti Connect Secure.</p>	
Enable User Record Synchronization	<p>Select this option to retain the bookmarks and individual preferences regardless of which system you log in to.</p>
Logical Auth Server Name	<p>Specify a logical authentication server name.</p>
Save Changes	<p>Click the button to save the changes made.</p>

Active Directory IPv6 Support

Active Directory server for authentication and authorization for AD mode auth server in ICS supports both IPv6 and IPv4 based backend Active Directory servers. If Active Directory server is configured with IPv6 only, then ICS is forced to use IPv6. If IPv6 is disabled in the backend server or in ICS, then ICS is forced to use IPv4. In case of a dual network in both the ICS and backend server, ICS would use both the protocols IPv6 and IPv4 for different authentication protocols like Kerberos, NTLM, etc.

ICS DNS server preferred mode settings do not apply to AD mode auth server since, internal third-party Samba library selects the available networks based on DNS resolution and other runtime protocol checks.

All features supported in IPv4 for Active Directory auth server are supported via IPv6 interface also.

Displaying the User Accounts Table

To display user accounts:

1. Select **Authentication > Auth. Servers**.
2. Click the link for the authentication server you want to manage.
3. Click the **Users** tab to display the user accounts table.

The user accounts table includes entries for the accounts that have been created. The Last Sign-in Statistic column shows the last successful sign-in date and time for each user, the user's IP address, and the agent or browser type and version.

4. Use the controls to search for users and manage user accounts:
 - To search for a specific user, enter a username in the Show users named field and click **Update**.

You can use an asterisk () as a wildcard, where * represents any number of zero or more characters. For example, to search for all usernames that contain the letters jo, enter *jo*. The search is case-sensitive. To display the entire list of accounts again, type * or delete the field's contents and click **Update**.*

- To limit the number of users displayed on the page, enter a number in the Show N users field and click **Update**.
- To terminate their user session and delete the account, select the check box next to the user account record and click **Delete**.

Troubleshooting the Active Directory Service

To troubleshoot the Active Directory Service:

1. Select **Authentication > Auth. Servers > AD Server name > Troubleshooting**.
2. Select the appropriate functions described in the following table.

The following table lists the Active Directory Server Troubleshooting Functions:

Function	Description
Basic Verification	Verifies whether the domain is properly joined and if the winbindd service is running. The number of winbindd processes is displayed, along with the ongoing CPU and memory usage for each process. For example, if user authentication is slow or fails randomly, use this function to check the number of winbindd processes and the CPU, memory and file descriptor usage. Select Restart AD Services to correct faulty processes.
Test User Authentication	Prompts for a username and password and attempts to log in. If successful, the groups the user belongs to are displayed. Only the regular password authentication is done.
Test User Password Change	Prompts for a username and the old and new password for a user and attempts to change the password on the AD server.
List Domain Info	Lists each domain and all trusted domains. Selecting a domain lists each Domain Controller for the domain, its IP address, and whether it is reachable. For example, if user authentication fails consistently and the domain is shown as successfully joined in the AD Server Settings page, the domain trust may be broken. Use this function to check the trusted domains. Also, if the domain join fails consistently or user authentication to a trusted domain fails consistently, the domain might not be reachable or the DNS configuration may be incorrect. Use this function to verify whether the domains and trusted domain are reachable.
Change Machine Password	Sends a request to the domain controller to change the machine password. A confirmation prompt is displayed to confirm the change.
Restart AD Services	Restarts the winbindd process, which may restore proper authentication, specifically during load and longevity scenarios. A confirmation prompt is displayed to confirm the restart (users cannot log in during a restart).

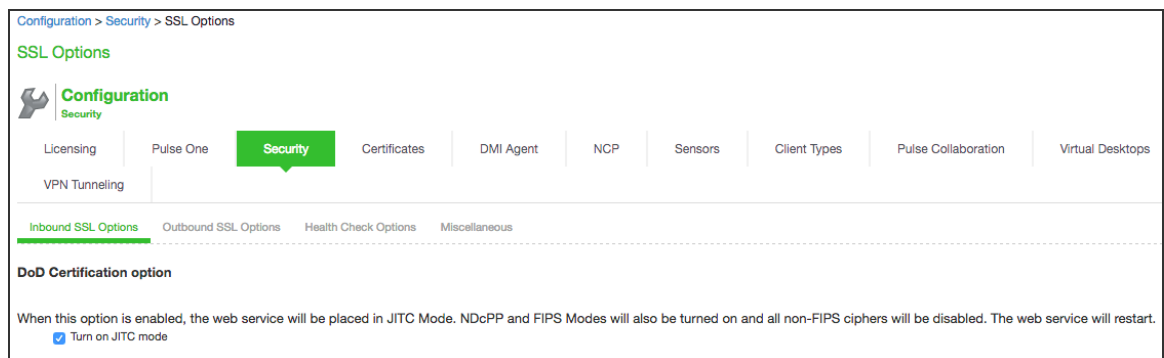
Function	Description
Reset Join	Reinitiates the domain join process. A confirmation prompt is displayed to confirm the reset and allows you to clear the Samba cache and keytab files before the reset. This is the same function shown on the AD server's Settings page and requires Active Directory administrator credentials. For example, if user group changes are not reflected in the user authentication, run this function with Clear Samba Cache enabled.
Samba Diagnostics Logs	Displays Diagnostic Logs page where you can download the Samba logs.
Load Output	Displays up to the last 500 lines of the troubleshooting output for the current session.
Save Output File	Saves all the troubleshooting messages for the current session.
Clear Output File	Erases all the troubleshooting messages saved in the output file (they cannot be retrieved).

JITC AAA Certification

Enabling JITC Mode

To enable the JITC Mode:

1. Navigate to **System > Configuration > Security > Inbound SSL Options**.
2. Click on **Turn on JITC mode** checkbox.



3. Once Turn on JITC mode is enabled, Turn on **NDcPP mode** and Turn on FIPS mode are also automatically enabled.

4. Click **Save Changes**.



For more details about the deployment of ICS in the JITC Mode, refer to the ICS/IPS NDcPP and JITC Certification Deployment Guide.

Important Factors to Consider

- **Password Strengthening:** When JITC is enabled, ICS does not allow an administrator to configure a password exactly same as previously configured 5 passwords. An error message is displayed in this case.
- **Notification for Unsuccessful Admin Login Attempts:** With JITC Mode on, ICS shows a banner with the count of unsuccessful login attempts. This includes any change in the admin status that would have happened since the last successful login. Upon clicking on the banner, the administrator is directed to the status page, which provides more details about status or configuration change since last log-in. These configuration changes are cleared before the next login so that admin can see different set of configuration changes, if anything happened from the last login.
- **Re-authentication of Admin Users:** ICS will force the administrator to re-authenticate with ICS whenever the following conditions occur:
 - Add Role
 - Delete Role
 - Modify the Role
 - Delete the Realm
 - Update the Realm
 - During DPE (Dynamic Policy Evaluation)
- **Configuration Change Notification:** For details about configuration changes and status information since last login, go to **System > Status > Admin Notification**.

Understanding Multidomain User Authentication

This topic provides an overview of multidomain user authentication with Active Directory and Windows NT.

Multi-Domain User Authentication Overview

The Ivanti access management framework allows for multidomain Active Directory and Windows NT authentication. The system authenticates users in the domain that you configure, users in child domains, and users in all domains trusted by the configured domain.

Users in the default domain can sign into the system using just their username, or the default domain and the username in the format default-domain\username.

When you enable trusted domain authentication, users in trusted or child domains can sign in using the name of the trusted or child domain and the username in the format trusted-domain\username. Note that enabling trusted domain authentication adds to the server response time.

Windows NT User Normalization

To support multidomain authentication, the Ivanti access management framework uses "normalized" Windows NT credentials when it contacts an Active Directory or Windows NT4 domain controller for authentication. Normalized Windows NT credentials include both the domain name and the username: domain\username. Regardless of how the user signs in (either using just a username or using the domain\username format), the access management framework always processes the username in domain\username format.

When a user signs in using only their username, the access management framework normalizes their Windows NT credentials as default-domain\username. Authentication succeeds only if the user is a member of the default domain.

When a user signs in using the domain\username format, the access management framework attempts to authenticate the user as a member of the domain the user specifies. Authentication succeeds only if the user-specified domain is a trusted or child domain of the default domain. If the user specifies an invalid or untrusted domain, authentication fails.

Two variables, <NTUser> and <NTDomain>, allow you to individually refer to Windows NT domain and username values. The system populates these two variables with the Windows NT domain and username information.

In role mapping rules, when you specify USER = john, the system treats this rule semantically as NTUser = john AND NTDomain = defaultdomain.

Kerberos Support

We recommend you configure the Ivanti access management framework to use the Kerberos authentication protocol with Windows domain controllers. When a user logs in to the system, the system performs Kerberos authentication and attempts to fetch the Kerberos realm name for the domain controller, as well as all child and trusted realms, using LDAP calls.

You can use Kerberos differently. You can specify the Kerberos realm name when configuring an Active Directory authentication server. We do not recommend this method for two reasons:

- You cannot specify more than one realm name. The system cannot then authenticate against child or trusted realms of the realm you specify.
- If you misspell the realm name, the system cannot authenticate users against the proper realm.

Windows NT4 Support

The Ivanti access management framework does not support Kerberos-based authentication in Windows NT4 domain controllers. The system uses NTLM with a backend Windows NT4 domain controller.

Understanding Active Directory and Windows NT Group Information Support

This topic describes support for polling group information from Active Directory and Windows NT servers.

Active Directory Group Information Overview

The Ivanti access management framework supports user group lookup in Domain Local, Domain Global, and Universal groups in the default domain, child domains, and all trusted domains. The system obtains group membership using one of three methods that have different capabilities:

- Group information in User's Security Context - Returns information about the user's Domain Global groups.
- Group information obtained using LDAP search calls - Returns information about the user's Domain Global groups and about the user's Universal groups if the access management framework queries the Global Catalog Server.
- Group information using native RPC calls - Returns information about the user's Domain Local Group.

With respect to role-mapping rules, the system attempts group lookup in the following order:

- Checks for all Domain Global groups using the user's security context.
- Performs an LDAP query to determine the user's group membership.
- Performs an RPC lookup to determine the user's Domain Local group membership.

Windows NT4 Group Information Overview

The Ivanti access management framework supports group lookup in the Domain Local and Domain Global groups created in the default domain, as well as all child and other trusted domains. The system obtains group membership using:

- Domain Global group information from the user's security context.
- Domain Local information using RPC calls.

In the Windows NT4 environment, the system does not use LDAP-based search calls.

Join Domain for Active Directory-based Authentication Server Without Using a Domain Admin Account

With Active Directory on Windows Server, the system can join domain (for an Active Directory based Authentication server) without using a domain administrator account. For more details refer to KB2624.

Using the Certificate Server

This topic describes integration with the certificate server.

Certificate Server Overview

This section describes support for using Ivanti Connect Secure with the certificate server.

Understanding the Certificate Server

The certificate server is a local server that allows user authentication based on the digital certificate presented by the user without any other user credentials.

If the certificate is secured through a hardware or a software token or through a password, the certificate server authentication is very useful. The certificate contains the full distinguished name (DN) and the system extracts the values from the DN and uses it for role mapping rules, authentication policies, and role restrictions.

Feature Support

The Ivanti access management framework supports the following certificate server features:

- Certificate directory services to retrieve user attributes in role mapping rules, authentication policies, and role restrictions.
- Load CA-created certificates on the system.
- Load multiple certificates from different CAs for use with different authentication realms.

Interoperability Requirements and Limitations

If you choose a certificate attribute with more than one value, the system uses the first matched value. For example, if you enter <certDN.OU> and the user has two values for the attribute (ou=management, ou=sales), the system uses the "management" value.

To use all values, add the SEP attribute to the variable. For example, if you enter <certDN.OUT SEP=":"> the system uses "management:sales".

Configuring Authentication with the Certificate Server

To configure authentication with the certificate server:

1. Select **Authentication > Auth. Servers**.
2. Select **Certificate Server** and click **New Server** to display the configuration page.
3. Complete the configuration as described in the following table.
4. Save the configuration.

The following table lists Certificate Server Settings

Settings	Guidelines
Name	Specify a name to identify the server within the system.

Settings	Guidelines
User Name Template	Specify a username template. Specify how the system should construct a username. You may use any combination of certificate variables contained in angle brackets and plain text. This value populates the <USER> and <USERNAME> session variables for use throughout the rest of the system configuration.
User Record Synchronization	This applies only to Ivanti Connect Secure.
Enable User Record Synchronization	Select this option to retain the bookmarks and individual preferences regardless of which system you log in to.
Logical Auth Server Name	Specify a logical authentication server name.

Displaying the User Accounts Table

To display user accounts, refer to [Displaying the User Accounts Table](#)

Using an LDAP Server

This topic describes integration with the LDAP server.

LDAP Server Overview

This section describes support for using Ivanti Connect Secure with the LDAP server.

Understanding LDAP Server

Lightweight Directory Access Protocol (LDAP) facilitates the access of online directory services. The Internet Engineering Task Force (IETF) designed and specified LDAP as a better way to make use of X.500 directories, having found the original Directory Access Protocol (DAP) too complex for average Internet clients to use. LDAP is a relatively simple protocol for updating and searching directories running over TCP/IP.

LDAP directory consists of a collection of attributes with a name, known as a distinguished name (DN). Each of the entry's attributes, known as a relative distinguished name (RDN), has a type and one or more values. The types are typically mnemonic strings, such as CN for common name. The valid values for each field depend on the types.

The full DN is constructed by stringing together RDNs from most specific to least specific, separated by commas, as shown in the following example:

```
cn=Bob_Employee, ou= account_mgr, o=sales, dc=Acme,dc=com.
```

LDAP Feature Support

Ivanti access management framework supports the following LDAP features:

- LDAP directory services to retrieve user attributes and group membership in role mapping rules
- Encrypted connections to the LDAP server using LDAP over SSL (LDAPS) or Start Transport Layer Security (TLS)
- Password management feature enabling users who access an LDAP server to manage their passwords using the policies defined on the LDAP server
- Fine-grained password policy (FGPP) for Active Directory 2008

Interoperability Requirements and Limitations

The following limitations apply to interoperability with LDAP:

- Backup LDAP servers must be the same version as the primary LDAP server. Also, we recommend that you specify the IP address of a backup LDAP server instead of its hostname, which might accelerate failover processing by eliminating the need to resolve the hostname to an IP address.

Configuring Authentication with an LDAP Server

The LDAP authentication configuration is enhanced in 9.1R3 to locate the nearest Microsoft domain controllers, which are spread across the globe, by resolving DNS SRV records.

To configure authentication with an LDAP server:

1. Select **Authentication > Auth. Servers**.
2. Select **LDAP Server** and click **New Server** to display the configuration page.
3. Complete the configuration as described in the following table.
4. Save the configuration.

The following table lists LDAP Server Settings

Settings	Guidelines
Name	Specify a name to identify the server within the system.
Enable Domain Name (enabled)	Select this option if you want to fetch a list of servers from the DNS server. Domain Name When you Enable Domain Name, specify the LDAP Domain name that can be mapped to domain controllers by DNS service.
Enable Domain Name (disabled)	Clear this option if you want to manually enter all the domain controllers host names.
	LDAP Server Specify the LDAP server name or the IP address.
	Backup LDAP Server1 (Optional) Specify the parameters for backup LDAP server1. Default port number: 389 (unencrypted connection). The specified backup LDAP server is used for failover processing. The authentication request is first routed to the primary LDAP server, and then to the specified backup servers if the primary server is unreachable.
	Backup LDAP Port1 Specify the parameters for backup LDAP port1.
	Backup LDAP Server2 (Optional) Specify the parameters for backup LDAP server2.
	Backup LDAP Port2 Specify the parameters for backup LDAP port2.
LDAP Port	Specify the LDAP port for the LDAP server. Default port number: 389 (unencrypted connection) Default port number: 636 (SSL connection)
LDAP Server Type	Select the backend LDAP server type from the following choices: Generic Active Directory

Settings	Guidelines
Connection	<p>Select one of the following options for the connection to the LDAP server:</p> <p>Unencrypted - The device sends the username and password to the LDAP Directory Service in cleartext.</p> <p>LDAPS - The device encrypts the data in the LDAP authentication session using the Secure Socket Layer (SSL) protocol before sending it to the LDAP Directory Service.</p> <p>Start TLS - The device allows both secure and plain requests against an LDAP server on a single connection.</p> <p>If you select LDAPS or Start TLS, the Validate Certificate option is displayed for the configured LDAP server(s) and its referral servers. Select this option if the SSL connection uses digital certificate security.</p> <p>If you enable validation for the referral servers, make sure your network DNS supports reverse lookup zone.</p> <p>If you want to verify the server certificates, the root CA and Intermediate CAs must be imported under trusted server CAs.</p>
Connection Timeout (seconds)	<p>Specify the time to wait for connection to the primary LDAP server, and then to each backup LDAP server.</p> <p>Default: 15 seconds</p>
Search Timeout (seconds)	<p>Specify the time to wait for search results from a connected LDAP server.</p>
Test Connection	<p>(Optional) To verify the connection between Pulse client and LDAP servers, click the Test Connection button.</p> <p>We recommend using the Test Connection function only after saving changes on the LDAP Server Configuration page.</p>
Authentication required?	

Settings	Guidelines
Authentication required to search LDAP	<p>Select this option to require authentication when performing search or password management operations.</p> <p>If you use Active Directory, you must select the Authentication required to search LDAP check box and provide the full DN and password of primary and backup administrator accounts that can reach Active Directory. You can enable password management on any LDAP server.</p> <p>This feature enables users who authenticate through an LDAP server to manage their passwords through the system using the policies defined on the LDAP server. To enable password management on any LDAP server, you must provide primary and backup administrator accounts (with write privileges to the directory) for the administrator DN and backup administrator DN.</p>
Admin DN	Specify the administrator DN for queries to the LDAP directory.
Password	Specify the password for the LDAP server.
Backup Admin DN	Specify the backup administrator DN for queries to the LDAP directory, as a fallback when primary Admin DN fails (due to account expiration). The interaction with LDAP directory stops when both primary and backup administrator accounts fail.
Backup Admin Password	Specify the backup administrator password for the LDAP server.
Finding user entries	
Base DN	Specify the base DN under which the users are located. For example, dc=sales,dc=acme, dc=com.
Filter	<p>Specify a unique variable that can be used to do a fine search in the tree. For example, samAccountname= <username> or cn= <username>.</p> <p>Include <username> in the filter to use the username entered on the sign-in page for the search.</p> <p>Specify a filter that returns 0 or 1 user DNs per user; the device uses the first DN returned if more than 1 DN is returned.</p>
Remove Domain from Windows users names?	

Settings	Guidelines
Strip domain from Windows username	Select this option to pass the username without the domain name to the LDAP server.
Determining group membership	
Base DN	Specify the base DN to search for user groups.
Filter	Specify a unique variable which can be used to do a fine search in the tree. For example, samAccountname= <username> or cn= <GROUPNAME>.
Member Attribute	Specify all the members of a static group.
Reverse group search	Select this option to start the search from the member instead of the group. This option is available only for Active Directory server types.
Query Attribute	Specify an LDAP query that returns the members of a dynamic group. For example, memberURL.
Nested Group Level	Specify how many levels within a group to search for the user. The higher the number, the longer the query time, so we recommend that you specify to perform the search no more than two levels deep.
Nested Group Search	Select one of the following options: Nested groups in Server Catalog - This option is faster because it can search within the implicit boundaries of the nested group. Search all nested groups - With this option, the device searches the Server Catalog first. If the device finds no match in the catalog, then it queries LDAP to determine if a group member is a subgroup.

Displaying the User Accounts Table

To display user accounts, refer to [Displaying the User Accounts Table](#)

Using the LDAP Password Management Feature

This topic describes support and limitations for LDAP password management.

LDAP Password Management Feature Overview

The password management feature enables users who access an LDAP server to manage their passwords through the Ivanti access management framework using the policies defined on the LDAP server. For example, if a user tries to sign in to the system with an LDAP password that is about to expire, the system notifies the user through the interface, and then passes the user's response back to the LDAP server without requiring the user to sign in to the LDAP server separately.

Users, administrators, and help desk administrators who work in environments where passwords have set expiration times may find the password management feature very helpful. If users are not informed that their passwords are about to expire, they can change them themselves through the system rather than call the help desk.

Once this feature is enabled, the system performs a series of queries to determine user account information, such as when the user's password was last set, whether the account is expired, and so on. The Ivanti access management framework does this by using its internal LDAP or Samba client. Many servers, such as Microsoft Active Directory offer an Administrative Console to configure account and password options.

LDAP-based password management works with the following types of LDAP servers:

- Microsoft Active Directory. For Active Directory, password policy attributes can be configured in the user entry container level or any organization level above the user container. If these attributes are configured at multiple levels, the level closest to the user node takes precedence. The password management feature is not supported on the Active Directory Global Catalog because password policy attributes are not fully populated in the Active Directory Global Catalog.
 - For Active Directory 2008, the Ivanti access management framework supports the Fine-Grained Password Policy (FGPP) configured in the AD user container.
- Generic LDAP servers such as OpenLDAP

The system relies on the back-end server to pinpoint the cause of error when a password change operation fails. However, although LDAP servers may report errors accurately to human operators, they do not always do so when communicating programmatically to systems. Therefore, reported errors might be generic or cryptic.

The system does not support customized password policies.

Enabling LDAP Password Management

To enable password management, you must first create an instance of the LDAP server. Next, you associate the LDAP server with the applicable realms. Finally, you select the enable password management feature at the realm level.

LDAP Password Management Support

The Ivanti access management framework supports password management with the following LDAP directories:

- Microsoft Active Directory/Windows NT

[Table](#) describes supported password management functions, their corresponding function names in the individual LDAP directories, and any additional relevant details. These functions must be set through the LDAP server itself before the system can pass the corresponding messages, functions, and restrictions to end users.

The Active Directory attribute names shown are specific to the Domain Security Policy object. Similar attributes for the corresponding functions are used for the Active Directory 2008 Fine-Grained Password Policy. Refer to Microsoft documentation for details.

When authenticating against a generic LDAP server, the system supports only authentication and allows users to change their passwords. Password management functions are not supported when the CHAP family protocols are used for authentication. All functions are available when the JUAC protocol is used for authentication (Policy Secure only).

The following table lists Supported Password Management Functions

Function	Active Directory	eDirectory
Authenticate user	unicodePwd	userPassword
Allow user to change password if enabled	Server tells us in bind response (uses ntSecurityDescriptor)	If passwordAllowChange == TRUE
Log out user after password change	Yes	Yes
Force password change at next login	If pwdLastSet == 0	If pwdMustChange == TRUE

Function	Active Directory	eDirectory
Expired password notification	userAccountControl == 0x80000	Check date/time value
Password expiration notification (in X days/hours)	if pwdLastSet - now() < maxPwdAge - 14 days (Read from domain attributes)	If now() - passwordExpirationTime < 14 days (The system displays warning if less than 14 days)
Disallow authentication if "account disabled/locked"	userAccountControl == 0x2 (Disabled) accountExpires userAccountControl == 0x10 (Locked) lockoutTime	Bind ErrorCode: 53 "Account Expired" Bind ErrorCode: 53 "Login Lockout"
Honor "password history"	"Server tells us in bind response"	Server tells us in bind response
Enforce "minimum password length"	"If set, the system displays message telling user minPwdLength"	If set, the system displays message telling user passwordMinimumLength
Disallow user from changing password too soon	If pwdLastSet - now() < minPwdAge, then we disallow	Server tells us in bind response
Honor "password complexity"	"If pwdProperties == 0x1, then enabled. Complexity means the new password does not contain username, first or last name, and must contain characters from 3 of the following 4 categories: English uppercase, English lowercase, Digits, and Non-alphabetic characters (ex. !, \$, %)"	Server tells us in bind response

Note the following expected behavior:

- The system displays a warning about password expiration only if the password is scheduled to expire in 14 days or less. The system displays the message during each sign-in attempt. The warning message contains the remaining number of days, hours, and minutes that the user has to change the password before it expires on the server. The default value is 14 days, but you can change it on the password configuration page of the admin console.

LDAP Password Management for Windows AD Versions

Note the following expected behavior:

- Changes on the Active Directory domain security policy can take 5 minutes or longer to propagate among Active Directory domain controllers. Additionally, this information does not propagate to the domain controller on which it was originally configured for the same time period. This issue is a limitation of Active Directory.
- When changing passwords in Active Directory using LDAP, the system automatically switches to LDAPS, even if LDAPS is not the configured LDAP method. To support LDAPS on the Active Directory server, you must install a valid SSL certificate into the server's personal certificate store. The certificate must be signed by a trusted CA, and the CN in the certificate's Subject field must contain the exact hostname of the Active Directory server, (for example: adsrv1.company.com). To install the certificate, select the Certificates Snap-In in the Microsoft Management Console (MMC).
- The Account Expires option in the User Account Properties tab only changes when the account expires, not when the password expires. Microsoft Active Directory calculates the password expiration using the Maximum Password Age and Password Last Set values retrieved from the User object and Fine-Grained Password Policy objects or the Domain Security Policy LDAP objects.
- The system displays a warning about password expiration only if the password is scheduled to expire in 14 days or less. The system displays the message during each sign-in attempt. The warning message contains the remaining number of days, hours, and minutes that the user has to change the password before it expires on the server. The default value is 14 days, but you can change it on the password configuration page of the admin console.

Troubleshooting LDAP Password Management

When you troubleshoot, provide any pertinent system logs, server logs, configuration information, and a TCP trace from the system. If you are using LDAPS, switch to the "Unencrypted" LDAP option LDAP server configuration while taking the LDAP TCP traces.

Using an MDM Server

This topic describes integration with the mobile device management (MDM) servers.

Understanding MDM Integration

MDM vendors provide enrollment and posture assessment services that prompt employees to enter data about their mobile devices. When the user installs the MDM application on the device and completes enrollment, the MDM pushes the device certificate to the device. After enrollment, the MDM maintains a database record that includes information about the enrollee-attributes related to device identity, user identity, and posture assessment against MDM policies.

The Ivanti access management framework MDM authentication server configuration determines includes details on how the system communicates with the MDM Web RESTful API service and how it derives the device identifier from the certificates presented by endpoints.

After you have configured the MDM authentication server, you can configure a realm that uses the MDM data for authorization, and you can use MDM device attributes in the role mapping rules that are the basis for your network access and resource access policies.

Feature Support

The Ivanti device access management framework supports integration with the following MDM solutions:

- Pulse Workspace
- AirWatch
- MobileIron
- Microsoft Intune

Configuring an MDM Server

The authentication server configuration is used by the system to communicate with the MDM. In the device access management framework, the MDM server is used as the device authorization server.

To configure the authentication server:

1. Select **Authentication > Auth Servers** to navigate to the authentication server configuration pages.

2. Select **MDM Server** and click **New Server** to display the configuration page.
3. Complete the configuration as described in the following table.
4. Save the configuration.

The following table lists Authentication Server Configuration Guidelines:

Settings	Guidelines
Name	Specify a name for the configuration.
Type	Select one of the following options: Pulse Workspace AirWatch MobileIron Microsoft Intune Ivanti Connect Secure has to be registered with Pulse One for using Pulse Workspace as an MDM auth server.
Server (Applicable to AirWatch and MobileIron)	
Server Url	Specify the URL for the MDM server. This is the URL the MDM has instructed you to use to access its RESTful Web API (also called a RESTful Web service). You must configure your firewalls to allow communication between these two nodes over port 443.
Viewer Url	Specify the URL for the MDM report viewer. This URL is used for links from the Active Users page to the MDM report viewer.
Request Timeout	Specify a timeout period (5-60 seconds) for queries to the MDM server. The default is 15 seconds.
Server (Applicable to Microsoft Intune)	
Tenant ID	Specify Azure AD Tenant ID.
Client ID	Specify Web application ID that has been registered in Azure AD.
Client Secret	Specify Secret key of the web application registered in azure AD.
Request Timeout	Specify a timeout period (5-60 seconds) for queries to the MDM server. The default is 15 seconds.
Administrator (Applicable to AirWatch and MobileIron)	

Settings	Guidelines
Username	Specify the username for an account that has privileges to access the MDM RESTful Web API.
Password	Specify the corresponding password.
Tenant Code	AirWatch only. Copy and paste the AirWatch API tenant code.
Device Identifier	
Device identity	<p>Ivanti Policy Secure only.</p> <p>Select an option on whether to require that the MDM certificate is presented by the endpoint when signing in:</p> <p>Require - Require that the device certificate pushed to client devices during enrollment be used at sign-in. If this option is selected, and the client device does not have a certificate, authorization fails. Use this option when you require endpoints to adhere to your certificate security requirements.</p> <p>Use Certificate if present - Use the certificate to derive the device ID if the certificate is presented at sign-in, but do not reject authentication if the certificate is not present. You can use this option in conjunction with a role mapping rule and a remediation VLAN to identify devices that have not perfected MDM enrollment.</p> <p>Always Use MAC address - In some cases, the MDM certificate might be configured without a device identifier. When the endpoint uses an 802.1x framework to authenticate, Ivanti Policy Secure can obtain the MAC address from the RADIUS return attribute callingStationID. The system can then use the MAC address as the device identifier.</p> <p>The Always Use MAC address option is not present in Ivanti Connect Secure. A device certificate is required to determine device identity.</p>
ID Template	<p>Construct a template to derive the device identifier from the certificate attributes. The template can contain textual characters as well as variables for substitution. The variables are the same as those used in role mapping custom expressions and policy conditions. Enclose variables in angle brackets like this <variable>.</p> <p>For example, suppose the certificate DN is: CN=<DEVICE_UDID>, uid=<USER_ID>, o=Company. With this configuration, the certificate could identify both the user and the device. In this example, the device ID template is <certDN.CN>.</p>

Settings	Guidelines
ID Type	<p>Select the device identifier type that matches the selection in the MDM SCEP certificate configuration:</p> <p>UUID - The device Universal Unique Identifier. This is the key device identifier supported by MobileIron MDM.</p> <p>Serial Number - The device serial number.</p> <p>UDID - The device Unique Device Identifier. This is the key device identifier supported by AirWatch MDM.</p> <p>IMEI - The device unique identifier. IMEI (15 decimal digits: 14 digits plus a check digit) or IMEISV (16 digits) includes information on the origin, model, and serial number of the device. This is the key device identifier supported by Microsoft Intune.</p>

Display the Active Users Page

The Active Users page lists data about current sessions, including access to realms that use the MDM server for device authorization.

To display the Active Users page, select **Systems > Active Users**.

The following figure shows the Active Users page for Ivanti Connect Secure:


Status > Active Users

Active Users

Activity Overview **Active Users** Cloud Secure Behavioral Analytics Virtual Desktop Sessions Devices Admin Notification

Show users named: Show users

Number of User Sessions: 1

	Primary User	Secondary User	Realm	Roles	Signed in	VPN Tunneling IP	VPN Tunnel Transport Mode	Device Details	Agent Type	Agent Version	Endpoint Security Status
<input type="checkbox"/>	admindb		Admin Users	Administrators	2021/2/4 20:11:20				Windows 8.1 Edge Chromium		Not Applicable
<input type="checkbox"/>	csaccess		Android_CloudSecure Realm	Android_CloudSecure_Role	2021/2/11 20:21:20	10.96.74.13	ESP		Android Pulse Secure		Not Applicable



Click the icon in the Device Details column to navigate to the MDM report viewer page for the device.

Using a RADIUS Server

This topic describes integration with the RADIUS server.

RADIUS Server Overview

This section describes support for using an external RADIUS server.

Understanding RADIUS Server

A Remote Authentication Dial-In User Service (RADIUS) server is a type of server that allows you to centralize authentication and accounting for users.

The following authentication schemes are supported:

- **Access-Request** - The user enters the username and password to request access to RADIUS server.
- **Access-Accept** - The user is authenticated.
- **Access-Reject** - The user is not authenticated and is prompted to reenter the username and password, or access is denied.
- **Access-Challenge** - A challenge is issued by the RADIUS server. The challenge collects additional data from the user.

Feature Support

Ivanti access management framework supports the following RADIUS features:

- RADIUS authentication.
- RADIUS attributes that can be used in role mapping.
- RADIUS directory services to retrieve user attributes in role-mapping rules.
- RADIUS accounting to track the services and the network resources used.
- RADIUS Disconnect messages. This feature is applicable for Ivanti Connect Secure.

Using Challenge Expressions

The Ivanti access management framework supports the RSA Authentication Manager using the RADIUS protocol and a SecurID token (available from Security Dynamics). If you use SecurID to authenticate users, they must supply a user ID and the concatenation of a PIN and a token value.

When you define a RADIUS server, the Ivanti access management framework allows administrators to use hard-coded (default) challenge expressions that support Defender 4.0 and some RADIUS server implementations (such as Steel-Belted RADIUS and RSA RADIUS) or to enter custom challenge expressions that allow the system to work with many different RADIUS implementations and new versions of the RADIUS server, such as Defender 5.0. The system looks for the response in the Access-Challenge packet from the server and issues an appropriate Next Token, New PIN, or Generic Passcode challenge to the user.

Using CASQUE Authentication

CASQUE authentication uses a token-based challenge/response authentication mechanism employing a CASQUE player installed on the client system. Once configured with CASQUE authentication, the RADIUS server issues a challenge with a response matching the custom challenge expression (:([0-9a-zA-Z/+ =]+):). The system then generates an intermediate page that automatically launches the CASQUE player installed on the user's system.

PassGo Defender

If you are using a PassGo Defender RADIUS server, the user sign-in process is as follows:

1. The user is prompted for and enters a username and password.
2. The username and encrypted password are sent over the network to the RADIUS server.
3. The RADIUS server sends a unique challenge string to the system. The system displays this challenge string to the user.
4. The user enters the challenge string in a Defender token and the token generates a response string.
5. The user enters the response string on the system and clicks Sign In.

Using RADIUS Attributes

The following table describes the RADIUS attributes that are supported in RADIUS role-mapping:

Attribute	Description
ARAP-Challenge-Response	Contains the response to the challenge of a dial-in client. Sent in an Access-Accept packet with Framed-Protocol of ARAP.

Attribute	Description
ARAP-Features	Includes password information that the network access server (NAS) must send to the user in an ARAP feature flags packet. Sent in an Access-Accept packet with Framed- Protocol of ARAP.
ARAP-Password	Appears in an Access-Request packet containing a Framed-Protocol of ARAP. Only one of User-Password, CHAP-Password, or ARAP-Password must be included in an Access-Request, or one or more EAP-Messages.
ARAP-Security	Identifies the ARAP security module to be used in an Access-Challenge packet.
ARAP-Security-Data	Contains the actual security module challenge or response, and is in Access-Challenge and Access-Request packets.
ARAP-Zone-Access	Indicates how to use the ARAP zone list for the user.
Access-Accept	Provides specific configuration information necessary to begin delivery of service to the user.
Access-Challenge	Sends the user a challenge requiring a response, and the RADIUS server must respond to the Access-Request by transmitting a packet with the Code field set to 11 (Access-Challenge). Access Challenge Response is not qualified over IPv6
Access-Reject	Transmits a packet with the Code field set to 3 (Access-Reject) if any value of the received Attributes is not acceptable.
Access-Request	Conveys information specifying user access to a specific NAS, and any special services requested for that user.
Accounting-Request	Conveys information used to provide accounting for a service provided to a user.
Accounting-Response	Acknowledges that the Accounting-Request has been received and recorded successfully.
Acct-Authentic	Indicates how the user was authenticated, whether by RADIUS, the NAS itself, or another remote authentication protocol.
Acct-Delay-Time	Indicates how many seconds the client has been trying to send this record.
Acct-Input-Gigawords	Indicates how many times the Acct-Input-Octets counter has wrapped around 2^{32} over the course of this service being provided.

Attribute	Description
Acct-Input-Octets	Indicates how many octets have been received from the port during the current session.
Acct-Input-Packets	Indicates how many packets have been received from the port during the session provided to a Framed User.
Acct-Interim-Interval	Indicates the number of seconds between each interim update in seconds for this specific session.
Acct-Link-Count	Indicates the count of links known to have been in a given multilink session at the time the accounting record is generated.
Acct-Multi-Session-Id	Indicates a unique Accounting ID to make it easy to link together multiple related sessions in a log file.
Acct-Output-Gigawords	Indicates how many times the Acct-Output-Octets counter has wrapped around 2^{32} during the current session.
Acct-Output-Octets	Indicates how many octets have been sent to the port during this session.
Acct-Output-Packets	Indicates how many packets have been sent to the port during this session to a Framed User.
Acct-Session-Id	Indicates a unique Accounting ID to make it easy to match start and stop records in a log file.
Acct-Session-Time	Indicates how many seconds the user has received service.
Acct-Status-Type	Indicates whether this Accounting-Request marks the beginning of the user service (Start) or the end (Stop).
Acct-Terminate-Cause	Indicates how the session was terminated.
Acct-Tunnel-Connection	Indicates the identifier assigned to the tunnel session.
Acct-Tunnel-Packets-Lost	Indicates the number of packets lost on a given link.
CHAP-Challenge	Contains the Challenge Handshake Authentication Protocol (CHAP) challenge sent by the NAS to a PPP CHAP user.
CHAP-Password	Indicates the response value provided by a PPP CHAP user in response to the challenge.

Attribute	Description
Callback-Id	Indicates the name of a location to be called, to be interpreted by the NAS.
Callback-Number	The dialing string to be used for callback.
Called-Station-Id	Allows the NAS to send the phone number that the user called, using Dialed Number Identification Service (DNIS) or similar technology.
Calling-Station-Id	Allows the NAS to send the phone number that the call came from, using Automatic Number Identification (ANI) or similar technology.
Class	Sent by the server to the client in an Access-Accept and then sent unmodified by the client to the accounting server as part of the Accounting-Request packet, if accounting is supported.
Configuration-Token	Used in large distributed authentication networks based on proxy.
Connect-Info	Sent from the NAS to indicate the nature of the user's connection.
Event-Timestamp	Records the time that this event occurred on the NAS, in seconds since January 1, 1970 00:00 UTC.
Filter-Id	Indicates the name of the filter list for this user.
Framed-AppleTalk-Link	Indicates the AppleTalk network number used for the serial link to the user, which is another AppleTalk router.
Framed-AppleTalk-Network	Indicates the AppleTalk Network number which the NAS can probe to allocate an AppleTalk node for the user.
Framed-AppleTalk-Zone	Indicates the AppleTalk Default Zone to be used for this user.
Framed-Compression	Indicates the compression protocol to be used for the link.
Framed-IP-Address	Indicates the address to be configured for the user.
Framed-IP-Netmask	Indicates the IP netmask to be configured for the user when the user is a router to a network.
Framed-IPv6-Pool	Contains the name of an assigned pool used to assign an IPv6 prefix for the user.
Framed-IPv6-Route	Indicates the routing information to be configured for the user on the NAS.
Framed-IPX-Network	Indicates the IPX Network number to be configured for the user.

Attribute	Description
Framed-MTU	Indicates the maximum transmission unit to be configured for the user, when it is not negotiated by some other means (such as PPP).
Framed-Pool	Indicates the name of an assigned address pool used to assign an address for the user.
Framed-Protocol	Indicates the framing to be used for framed access.
Framed-Route	Indicates the routing information to be configured for the user on the NAS.
Framed-Routing	Indicates the routing method for the user, when the user is a router to a network.
Idle-Timeout	Sets the maximum number of consecutive seconds of idle connection allowed to the user before termination of the session or prompt.
Keep-Alives	Uses SNMP instead of keepalives.
Login-IP-Host	Indicates the system with which to connect the user when the Login-Service Attribute is included.
Login-IPv6-Host	Indicates the system with which to connect the user when the Login-Service Attribute is included.
Login-LAT-Group	Contains a string identifying the LAT group codes that this user is authorized to use.
Login-LAT-Node	Indicates the node with which the user is to be automatically connected by LAT.
Login-LAT-Port	Indicates the port with which the user is to be connected by LAT.
Login-LAT-Service	Indicates the system with which the user is to be connected by LAT.
Login-Service	Indicates the service to use to connect the user to the login host.
Login-TCP-Port	Indicates the TCP port with which the user is to be connected when the Login-Service Attribute is also present.
MS-ARAP-Challenge	Only present in an Access-Request packet containing a Framed-Protocol Attribute with the value 3 (ARAP).
MS-ARAP-Password-Change-Reason	Indicates the reason for a server-initiated password change.
MS-Acct-Auth-Type	Represents the method used to authenticate the dial-up user.

Attribute	Description
MS-Acct-EAP-Type	Represents the Extensible Authentication Protocol (EAP) type used to authenticate the dial-up user.
MS-BAP-Usage	Describes whether the use of BAP is allowed, disallowed, or required on new multilink calls.
MS-CHAP-CPW-1	Allows the user to change password if it has expired.
MS-CHAP-CPW-2	Allows the user to change password if it has expired.
MS-CHAP-Challenge	Contains the challenge sent by a NAS to a MS-CHAP user.
MS-CHAP-Domain	Indicates the Windows NT domain in which the user was authenticated.
MS-CHAP-Error	Contains error data related to the preceding MS-CHAP exchange.
MS-CHAP-LM-Enc-PW	Contains the new Windows NT password encrypted with the old LAN Manager password hash.
MS-CHAP-MPPE-Keys	Contains two session keys for use by the Microsoft Point-to-Point Encryption (MPPE).
MS-CHAP-NT-Enc-PW	Contains the new Windows NT password encrypted with the old Windows NT password hash.
MS-CHAP-Response	Contains the response value provided by a PPP MS-CHAP user in response to the challenge.
MS-CHAP2-CPW	Allows the user to change password if it has expired.
MS-CHAP2-Response	Contains the response value provided by an MS-CHAP-V2 peer in response to the challenge.
MS-CHAP2-Success	Contains a 42-octet authenticator response string.
MS-Filter	Transmits traffic filters.
MS-Link-Drop-Time-Limit	Indicates the length of time (in seconds) that a link must be underutilized before it is dropped.
MS-Link-Utilization-Threshold	Represents the percentage of available bandwidth utilization below which the link must fall before the link is eligible for termination.
MS-MPPE-Encryption-Policy	Signifies whether the use of encryption is allowed or required.

Attribute	Description
MS-MPPE-Encryption-Types	Signifies the types of encryption available for use with MPPE.
MS-MPPE-Recv-Key	Contains a session key for use by the MPPE.
MS-MPPE-Send-Key	Contains a session key for use by the MPPE.
MS-New-ARAP-Password	Transmits the new ARAP password during an ARAP password change operation.
MS-Old-ARAP-Password	Transmits the old ARAP password during an ARAP password change operation.
MS-Primary-DNS-Server	Indicates the address of the primary domain name server (DNS) server to be used by the PPP peer.
MS-Primary-NBNS-Server	Indicates the address of the primary NetBIOS name server (NBNS) server to be used by the PPP peer.
MS-RAS-Vendor	Indicates the manufacturer of the RADIUS client machine.
MS-RAS-Version	Indicates the version of the RADIUS client software.
MS-Secondary-DNS-Server	Indicates the address of the secondary DNS server to be used by the PPP peer.
MS-Secondary-NBNS-Server	Indicates the address of the secondary DNS server to be used by the PPP peer.
Message-Authenticator	Signs Access-Requests to prevent spoofing Access-Requests using CHAP, ARAP, or EAP authentication methods.
NAS-IP-Address	Indicates the identifying IP address of the NAS that is requesting authentication of the user, and must be unique to the NAS within the scope of the RADIUS server.
NAS-IPv6-Address	Indicates the identifying IPv6 Address of the NAS that is requesting authentication of the user, and must be unique to the NAS within the scope of the RADIUS server.
NAS-Identifier	Contains a string identifying the NAS originating the Access-Request.
NAS-Port	Indicates the physical port number of the NAS that is authenticating the user.

Attribute	Description
NAS-Port-Id	Contains a text string that identifies the port of the NAS that is authenticating the user.
NAS-Port-Type	Indicates the type of the physical port of the NAS that is authenticating the user.
Password-Retry	Indicates how many authentication attempts a user is allowed to attempt before being disconnected.
Port-Limit	Sets the maximum number of ports to be provided to the user by the NAS.
Prompt	Indicates to the NAS whether it should echo the user's response as it is entered, or not echo it.
Proxy-State	Indicates that a proxy server can send this attribute to another server when forwarding an Access-Request. The attribute must be returned unmodified in the Access-Accept, Access-Reject or Access-Challenge.
Reply-Message	Indicates that the text that can be displayed to the user.
Service-Type	Indicates the type of service the user has requested, or the type of service to be provided.
Session-Timeout	Sets the maximum number of seconds of service to be provided to the user before termination of the session or prompt.
State	Indicates that the packet must have only zero or one State Attribute. Usage of the State Attribute is implementation dependent.
Telephone-number	Using the Calling-Station-Id and Called-Station-Id RADIUS attributes, authorization and subsequent tunnel attributes can be based on the phone number originating the call, or the number being called.
Termination-Action	Indicates the action the NAS should take when the specified service is completed.
Tunnel-Assignment-ID	Indicates to the tunnel initiator the particular tunnel to which a session is to be assigned.
Tunnel-Client-Auth-ID	Specifies the name used by the tunnel initiator during the authentication phase of tunnel establishment.
Tunnel-Client-Endpoint	Contains the address of the initiator end of the tunnel.

Attribute	Description
Tunnel-Link-Reject	Indicates the rejection of the establishment of a new link in an existing tunnel.
Tunnel-Link-Start	Marks the creation of a tunnel link.
Tunnel-Link-Stop	Marks the destruction of a tunnel link.
Tunnel-Medium-Type	Indicates the transport medium to use when creating a tunnel for those protocols (such as L2TP) that can operate over multiple transports.
Tunnel-Medium-Type	Indicates the transport medium to use when creating a tunnel for those protocols (such as L2TP) that can operate over multiple transports.
Tunnel-Password	Specifies a password used to access a remote server.
Tunnel-Preference	Indicates that if RADIUS server returns more than one set of tunneling attributes to the tunnel initiator, you should include this attribute in each set to indicate the relative preference assigned to each tunnel.
Tunnel-Private-Group-ID	Indicates the group ID for a particular tunneled session.
Tunnel-Reject	Marks the rejection of the establishment of a tunnel with another node.
Tunnel-Server-Auth-ID	Specifies the name used by the tunnel terminator during the authentication phase of tunnel establishment.
Tunnel-Server-Endpoint	Indicates the address of the server end of the tunnel.
Tunnel-Start	Marks the establishment of a tunnel with another node.
Tunnel-Stop	Marks the destruction of a tunnel to or from another node.
Tunnel-Type	Indicates the tunneling protocol(s) to be used (in the case of a tunnel initiator) or the tunneling protocol in use (in the case of a tunnel terminator).
User-Name	Indicates the name of the user to be authenticated.
User-Password	Indicates the password of the user to be authenticated, or the user's input following an Access-Challenge.

Understanding RADIUS Accounting

You can configure the device to send session start and stop messages to a RADIUS accounting server. The device sends a user-session start message after the user successfully signs in and the device maps to a role.

Whenever a user session is terminated, the device sends a user-session stop message to the accounting server. A user session is terminated whenever the user:

- Manually signs out
- Times out because of either inactivity or exceeding the maximum session length
- Is denied access because of Host Checker role-level restrictions
- Is manually forced out by an administrator as a result of dynamic policy evaluation



If users are signed into a device cluster, the RADIUS accounting messages might show the users signing in to one node and signing out of another.

The following table describes the attributes that are common to start and stop messages. The next table describes the attributes that are unique to start messages.

The following table lists the Attributes Common to Start and Stop Messages:

Attribute	Description
User-Name (1)	Specifies the string that the device administrator specifies during RADIUS server configuration.
NAS-IP-Address (4)	Specifies the device's IPv4 address.
NAS-IPV6-Address	Specifies the device's IPv6 address.
NAS-Port (5)	The device sets this attribute to 0 if the user signed in using an internal port, or 1 if an external port is used.
Framed-IP-Address (8)	Specifies the user's source IPv4 address.
Framed-IPv6-Address	Specifies the user's source IPv6 address.
NAS-Identifier (32)	Specifies the configured name for the device client under the RADIUS server configuration.

Attribute	Description
Acct-Status-Type (40)	The device sets this attribute to 1 for a start message, or 2 for a stop message in a user-session or a sub-session.
Acct-Session-Id (44)	Specifies the unique accounting ID that matches start and stop messages corresponding to a user-session or to a sub-session.
Acct-Multi-Session-Id (50)	Specifies the unique accounting ID that you can use to link together multiple related sessions. Each linked session must have a unique Acct-Session-Id and the same Acct-Multi-Session-Id.
Acct-Link-Count (51)	Specifies the count of links in a multilink session at the time the system generates the accounting record.

The following table lists Start Attributes:

Attribute	Description
Acct-Authentic (45)	The device sets this attribute to: RADIUS - if the user is authenticated to a RADIUS server. Local - if the user is authenticated to a local authentication server. Remote - if the user is authenticated through any other RADIUS server.

The following table lists Stop Attributes:

Attribute	Description
Acct-Session-Time (46)	Specifies the duration of the user-session or the sub-session.
Acct-Terminate-Cause (49)	The device uses one of the following values to specify the event that caused the termination of a user session or a sub-session: User Request (1) - User manually signs out. Idle Timeout (4) - User is idle and times out. Session Timeout (5) - User's maximum session times out. Admin Reset (6) - User is forced out from active users page.

Interoperability Requirements and Limitations

You must configure the third-party RADIUS server to communicate with the Ivanti access management framework.

On the RADIUS server, configure the following settings:

- Hostname.
- Network IP address.
- Client type, if applicable. If this option is available, select Single Transaction Server or its equivalent.
- Type of encryption for authenticating client communication. This choice should correspond to the client type.
- Shared secret.

The following are the requirements and limitations for Interim update feature:

- If you want a server to receive interim accounting messages, you can statically configure an interim value on the client, in which case, the locally configured value overrides any value that might be included in the RADIUS Access-Accept message.
- The octet count reported in the accounting messages is the cumulative total since the beginning of the user session.
- The interim update byte count is only supported based on a user session, not on SAM or NC sessions.

Configuring Authentication with a RADIUS Server

To configure authentication with the RADIUS server:

1. Select **Authentication > Auth. Servers**.
2. Select **RADIUS Server** and click **New Server** to display the configuration page.
3. Complete the configuration as described in [Table](#).
4. Save the configuration.

The following table lists RADIUS Server Settings:

Settings	Guidelines
Name	Specify a name to identify the server within the system.

Settings	Guidelines
NAS-Identifier	<p>Specify the name that identifies the Network Access Server (NAS) client to the RADIUS server.</p> <p>If you do not specify the NAS identifier, the value specified in the Hostname field on the System > Network > Overview page of the administrator console is used.</p> <p>If you use the RADIUS proxy feature, the NAS-Identifier field is not used. Proxy passes on the entire RADIUS packet including the NAS identifier from the client.</p>
Primary Server	
Radius Server	Specify the name or IP address of the RADIUS server.
Authentication Port	<p>Specify the authentication port value for the RADIUS server.</p> <p>Default port number: 1812, 1645 (legacy servers)</p>
NAS-IP-Address	<p>Specify the NAS IP address.</p> <p>If you leave this field empty, the internal IP address is passed to RADIUS requests.</p> <p>You can also fill this field with IPv6 address.</p> <p>If you configure the NAS IP address, then the system passes the value regardless of which cluster node sends the requests.</p> <p>If you use the RADIUS proxy feature, this field is not used.</p> <p>Proxy passes on the entire RADIUS packet including the NAS IP address from the client.</p>
Timeout (seconds)	Specify the interval of time to wait for a response from the RADIUS server before timing out the connection.
Retries	Specify the number of times to try to make a connection after the first attempt fails.
Users authenticate using tokens or one-time passwords.	<p>Select this option to prompt the user for a token instead of a password. For example, you can use this option to dynamically prompt for a password or token based on sign-in policies by configuring two instances of the same authentication server. You can use one instance for wireless users with this option enabled and that prompts the user for a token, and another instance for wired users with this option disabled and that prompts the user for a password.</p> <p>If you are using RADIUS proxy feature, this option is not used.</p>
Backup Server (required only if Backup server exists)	

Settings	Guidelines
Radius Server	<p>Specify the secondary RADIUS server.</p> <p>The authentication request is first routed to the primary RADIUS server, then to the specified backup server if the primary server is unreachable. Accounting messages are sent to the RADIUS server by each cluster node without consolidation.</p> <p>RADIUS accounting follows these assumptions:</p> <p>If the cluster is active/passive, all users are connected to one node at a time.</p> <p>If the cluster is active/active and does not use a balancer, users are connected to different nodes but are static.</p> <p>If the cluster is active/active and uses a balancer, the balancer usually enforces a persistent source IP. In this case, users are always connected to the same node.</p> <p>RADIUS does not support load balancing.</p>
Authentication Port	Specify the authentication port.
Shared Secret	Specify the shared secret.
Accounting Port	Specify the accounting port.
Radius Accounting	
User-Name	<p>Specify the user information to the RADIUS accounting server.</p> <p>You can enter any of the applicable session variables. Applicable variables include those that are set the time after the user signs in and maps to a role.</p> <p>The default variables for this field are as follows:</p> <p>USER: Logs the username to the accounting server.</p> <p>REALM: Logs the realm to the accounting server.</p> <p>ROLE SEP=",";: Logs the list of comma-separated roles assigned to the user.</p> <p>ROLE: Logs the role to the accounting server.</p> <p>If you assign the user to more than one role, the system separates them with commas.</p> <p>From 9.1R15, including password variable is not supported.</p>

Settings	Guidelines
Interim Update Interval (minutes)	<p>Select this option to achieve more precise billing for long-lived session clients and during network failure.</p> <p>If you are using the RADIUS proxy feature, the fields in this section are not used.</p> <p>The minimum interim update interval is 15 minutes. The data statistics (bytes in and bytes out) for RADIUS accounting might not be sent for a J-SAM/W-SAM/NC session if the session is less than 30 seconds long and the applications keep the connections open all the time.</p>
Send Interim Updates for sub sessions created inside parent sessions	<p>Enable this checkbox to send interim updates for sub sessions (child sessions) created inside parent sessions.</p>
Use VPN Tunnel assigned IP Address for FRAMED-IP-ADDRESS/FRAMED-IPV6-ADDRESS attribute value in RADIUS Accounting	<p>Select the Use NC assigned IP Address for FRAMED-IP-ADDRESS/FRAMED-IPV6-ADDRESS attribute value in Radius Accounting check box to use the IP address returned from Ivanti Connect Secure for the Framed-IP-Address attribute. Two IP addresses are recorded: one prior to authenticating with Ivanti Connect Secure, and one returned by VPN Tunneling after authentication. Select this option to use the VPN Tunneling IP address for the FRAMED-IP-ADDRESS/FRAMED-IPV6-ADDRESS attribute instead of the pre-authenticated (original) IP address.</p> <p>Framed IPv6 addresses based attribute fetching and parsing: NAS-IPv6-Address Login-IPv6-Host</p>
<p>Radius Disconnect This feature is applicable for Ivanti Connect Secure</p>	
Enable processing of Radius Disconnect Requests	<p>Select this option to process Radius Disconnect Requests. The Radius Disconnect requests received from the backend Radius server will terminate sessions that match the attributes in the request.</p> <p>You must not configure multiple RADIUS authentication servers with the same backend server details. Radius Disconnect over IPv6 is not qualified.</p> <p>The Radius attributes that are used for session identification are: Framed-IP-Address (for sessions with VPN Tunnel only) Acct-Session-Session-Id Acct-Multi-Session-Id User-Name</p>

Settings	Guidelines
Next Token	Specify the appropriate Next Token.
New PIN	Specify the New PIN.
Generic Login	Specify the Generic Login challenge to the user.
Custom Radius Rules This feature is applicable for Ivanti Connect Secure	

Settings	Guidelines
	<p>(Optional) Click New Radius Rule to add a custom challenge rule that determines the action to take for an incoming packet. When a user enters his or her username and password, the initial authorization request is sent to the server. The server may respond with a Challenge or Reject packet. In the Add Custom Radius Challenge Rule window, you select the packet type (Challenge or Reject) and then specify what action to take. For example, you can show a login page with a specific error message to the user, or automatically send an ACCESS-REQUEST packet back to the server. To create a custom challenge rule:</p> <ol style="list-style-type: none"> Select the incoming packet type: <ul style="list-style-type: none"> Access Challenge - sent by the RADIUS server requesting more information in order to allow access Access Reject - sent by the RADIUS server rejecting access Specify an expression to evaluate, based on the Radius attribute, and click Add. If you specify more than one expression, the expressions are "ANDed" together. To remove an expression, click the delete icon next to the expression. Choose the action to take by selecting one of the following radio buttons: <ul style="list-style-type: none"> show NEW PIN page - user must enter a new PIN for the token show NEXT TOKEN page - user must enter the next tokencode show GENERIC LOGIN page - display an additional page to the user in response to an Access Challenge sent by the server. Sometimes a Radius server returns a Challenge packet and requires the user to enter additional information to continue the login process. For example, a server receives the initial username and password and sends an SMS message to the user's mobile phone with a one-time password (OTP). The user enters the OTP in the generic login page. show user login page with error - display the standard login page with an embedded error message. This option lets you bypass the standard message string sent by Ivanti Connect Secure and display a custom error message to the user. Enter your custom message in the Error Message text box. There is no maximum character limit for this message. send ACCESS REQUEST with additional attributes - send an ACCESS-REQUEST packet with the specified attribute/value pair(s). Select an attribute, enter its value and click Add. To delete an attribute, click the delete icon next to the attribute/value pair. <p>You must set User-Password to <PASSWORD> otherwise an "Invalid username or password" message appears.</p> <p>Copyright © 2023 Ivanti, Inc. All Rights Reserved. Privacy and Legal</p>

Settings	Guidelines
----------	------------

Displaying the User Accounts Table

To display user accounts, refer to the steps found in the [Displaying the User Accounts Table](#) section.

Using an ACE Server

This topic describes integration with an ACE Server (now named RSA Authentication Manager).

RSA Authentication Manager Overview

This section describes support for using Ivanti Connect Secure with an ACE Server (now named RSA Authentication Manager).

Understanding RSA Authentication Manager

RSA Authentication Manager (formerly known as ACE/Server) is an authentication and authorization server that allows user authentication based on credentials from the RSA SecurID® product from RSA Security Inc.

When you use RSA Authentication Manager as the authentication and authorization service for your Ivanti access management framework, users can sign in to Ivanti Connect Secure using the same username and password stored in the backend server.

The following table describes RSA SecurID hardware token and software token user sign-in methods:

Method	Action
Using a hardware token and the standard system sign-in page	The user browses to the standard system sign-in page, and then enters the username and password (consisting of the concatenation of the PIN and the RSA SecurID hardware token's current value). The system then forwards the user's credentials to the authentication server.
Using a software token and the custom SoftID system sign-in page	The user browses to the SoftID custom sign-in page. Then, using the SoftID plug-in, the user enters the username and PIN. The SoftID plug-in generates a passphrase by concatenating the user's PIN and token and passes the passphrase to the authentication server.

If the RSA Authentication Manager positively authenticates the user, the user gains access to the system. Otherwise, the RSA Authentication Manager:

- Denies the user access to the system.
- Prompts the user to generate a new PIN (New PIN mode) if the user is signing in to the system for the first time. Users see different prompts depending on the method they use to sign in.
- If the user signs in using the SoftID plug-in, then the RSA prompts the user to create a new pin; otherwise Ivanti Connect Secure prompts the user to create a new PIN.
- Prompts the user to enter the next token (Next Token mode) if the token entered by the user is out of sync with the token expected by RSA Authentication Manager. Next Token mode is transparent to users signing in using a SoftID token. The RSA SecurID software passes the token through the system to RSA Authentication Manager without user interaction.
- Redirects the user to the standard system sign-in page (SoftID only) if the user tries to sign-in to the RSA SecurID Authentication page on a computer that does not have the SecurID software installed.

Feature Support

Ivanti access management framework supports the following RSA Authentication Manager features:

- **New PIN mode**
- **Next-token mode**
- **Data Encryption Standard (DES)/ Secure Dial-In (SDI) encryption**
- **Advanced Encryption Standard (AES) encryption**
- **Slave Authentication Manager support**
- **Name locking**
- **Clustering**

Interoperability Requirements and Limitations

The following limitations apply when defining and monitoring an RSA Authentication Manager instance:

- You can only add one RSA Authentication Manager configuration to the system, but you can use that configuration to authenticate any number of realms.
- You cannot customize the load balancing algorithm.

- When you enter the New PIN or Next Token mode, enter the required information within three minutes. Otherwise, the system cancels the transaction and notifies the user to reenter the credentials.
- The system can handle a maximum of 200 RSA Authentication Manager transactions at any given time. A transaction only lasts as long as is required to authenticate against the RSA Authentication Manager.

For example, when a user signs into the system, the RSA Authentication Manager transaction is initiated when the user submits the request for authentication and ends once the RSA Authentication Manager has finished processing the request. The user may then keep his or her session open, even though the RSA Authentication Manager transaction is closed.

Configuring Authentication with RSA Authentication Manager

To configure authentication with an ACE server:

1. Select **Authentication > Auth. Servers**.
2. Select **ACE Server** and click **New Server** to display the configuration page. Complete the configuration as described in the following table.
3. Save the configuration.

The following table lists the ACE Server Settings:

Settings	Guidelines
Name	Specify a name to identify the server within the system.
ACE Port	Specify the default port of the authentication server. If no port is specified in the sdconf.rec file, the default port is used.
Configuration File	
Current config file	Specify the RSA Authentication Manager configuration file. You must update this file on the device anytime you make changes to the source file.
Imported on	Display the date on which the config file is imported.
Import new config file	Use the Choose File button to upload the sdconf.rec configuration file.
Node Verification File	

Settings	Guidelines
Node	Save the configuration to redisplay the configuration page. The updated page includes a section that lists a timestamp for the negotiation of the node secret between the system and the backend RSA server. The negotiation and verification automatically occur after first successful login. Do not expect entries in the table until at least one user has authenticated successfully.
User Record Synchronization	This feature is available only on Ivanti Connect Secure.
Enable User Record Synchronization	Select this option to retain the bookmarks and individual preferences regardless of which system you log in to.
Logical Auth Server Name	Specify a logical authentication server name.

Enabling RSA Risk Based Authentication (RBA) Support with ICS Cluster

RSA SecurID Risk-Based Authentication is a token less, multi-factor enterprise authentication solution. ICS integration with Risk based authentication works with the usage of custom sign in pages.

1. Open the **ICS login** page.
2. ICS immediately delegates authentication to RSA server by redirecting the user **RSA Authentication Manager (AM)** server to authenticate.
3. User is now prompted for step-up authentication based on the risk score. For example: The user is challenged to answer enter additional security questions if the user logs in from a different endpoint.
4. Once successfully authenticated to RSA AM, the user is redirected back to ICS with a one-time token key, validated by ICS.
5. Each agent in RSA AM is linked to an agent ID in the integration file. Download this file from RSA AM and add to custom sign-in page package.
6. In case of cluster (for example 2 node cluster) two integration files (node1.js and node2.js) are required in the custom sign-in page package and it can be used in LoginPage.html.

For Example:

If the cluster node names are "node1" & "node2", add the similar lines inside the body (before the end) of LoginPage.html.

```
<% IF loginNode == "node1" %>

<script src='<% Home %>/node1.js' type="text/javascript"></script>

<% ELSE %>

<script src='<% Home %>/node2.js' type="text/javascript"></script>

<% END %>

<script>window.onload=redirectToIdP;</script>
```

7. In case of standalone ICS, the above conditional check with loginNode is not required. If the integration file name is am_integration.js, then add the integration file as part of custom sign-in page package and the below changes in LoginPage.html in Custom signin page would be sufficient.

```
<script src='<% Home %>/am_integration.js'
type="text/javascript"></script>

<script>window.onload=redirectToIdP;</script>
```



Also, all the related LoginPage-*.html (like LoginPage-ipad.html in Custom) needs similar changes to reflect the RBA login experience for browser-based login from different devices.

Displaying the User Accounts Table

To display user accounts, refer to the steps found in the [Displaying the User Accounts Table](#) section.

Using the SAML Server

This topic describes the local SAML authentication server.

SAML Server Overview

This section describes support for using the local Ivanti Connect Secure SAML authentication server.

Understanding SAML

SAML is an XML-based framework for communicating user authentication, entitlement, and attribute information. The standard defines the XML-based assertions, protocols, bindings, and profiles used in communication between SAML entities. SAML is used primarily to implement Web browser single sign-on (SSO). SAML enables businesses to leverage an identity-based security system like Ivanti Connect Secure to enforce secure access to web sites and other resources without prompting the user with more than one authentication challenge.

For complete details on the SAML standard, see the OASIS web site:

http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security

SAML Feature Support

When deployed as SAML service provider, Ivanti Connect Secure runs a local SAML server that relies on the SAML identity provider authentication and attribute assertions when users attempt to sign in to Ivanti Connect Secure. Note that authentication is only part of the Ivanti Connect Secure security system. The access management framework determines access to the system and protected resources.

Ivanti Connect Secure supports :

- HTTP Redirect binding for sending AuthnRequests
- HTTP Redirect binding for sending/receiving SingleLogout requests/responses
- HTTP POST and HTTP Artifact bindings for receiving SAML responses
- RequestedAuthnContext context class specifications

Interoperability Requirements and Limitations

Before you begin:

- Check to see whether the SAML identity provider implements SAML 2.0 or SAML 1.1.
- Check to see whether the SAML identity provider uses HTTP POST or HTTP Artifact bindings for SAML assertions.
- Check to see whether the SAML identity provider has published a SAML metadata file that defines its configuration. If the SAML identity provider metadata file is available, configuration is simpler and less prone to error.

- Complete the system-wide SAML settings if you have not already done so. Select **System > Configuration > SAML > Settings**. For details, see [Configuring Global SAML Settings](#).
- Add metadata for the SAML identity provider to the metadata provider list if you have not already done so. Select **System > Configuration > SAML**. For details, see [Managing SAML Metadata Files](#)

The sign-in URL for which a session needs to be established for Ivanti Connect Secure as a service provider is identified by the RelayState parameter (HTTP URL parameter for artifact and HTML form parameter for POST.) In a service provider initiated case, the system populates RelayState as an HTTP URL parameter while sending AuthnRequest. In the IdP-Initiated scenario (Ivanti Connect Secure is a service provider and there is a third-party IdP), the IdP must be configured to set the appropriate Sign-in URL of Ivanti Connect Secure in the RelayState parameter of the HTML form containing the SAML response. For more information, see the SAML 2.0 specification.

Configuring Authentication with the SAML Server

To configure the SAML server:

1. Select **Authentication > Auth. Servers**.
2. Select SAML Server and click New Server to display the configuration page.
3. Complete the configuration as described in the following table.
4. Save the configuration.

The following table lists SAML Service Provider Profile:

Settings	Guidelines
Name	Specify a name to identify the server instance.
Settings	
SAML Version	Select 2.0 or 1.1, depending on the SAML version used by the SAML IdP.
Override Global Host FQDN	Select to override the Host FQDN for the SAML server. Host FQDN is used to update the Unique SAML Identifier and ACS URL of the SAML Authentication Server.
Connect Secure Entity Id	This value is prepopulated. It is generated by the system, based on the value for the Host FQDN for SAML setting on the System > Configuration > SAML > Settings page.

Settings	Guidelines
Configuration Mode	<p>Select Manual or Metadata. If a metadata file or location is available from the SAML identity provider, use the metadata option to make configuration simpler and less prone to error. To upload or set the location for the published metadata file, select System > Configuration > SAML and click the New Metadata Provider button.</p>
Identity Provider Entity ID	<p>The identity provider entity ID is sent as the Issuer value in the assertion generated by the SAML identity provider.</p> <p>If you use the metadata option, this setting can be completed by selecting the identity provider entity ID from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page.</p> <p>If you complete this setting manually, specify the Issuer value in assertions generated by the SAML identity provider. Typically, you ask the SAML identity provider administrator for this setting.</p>
Identity Provider Single Sign On Service URL	<p>The identity provider SSO service URL is a URL provisioned by the SAML identity provider. The setting is required to support service-provider-initiated SSO. If missing, the system cannot successfully redirect the user request.</p> <p>If you use the metadata option, this setting can be completed by selecting the SSO service URL from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page.</p> <p>If you complete this setting manually, ask the SAML identity provider administrator for this setting.</p>

Settings	Guidelines
User Name Template	<p>Specify how the system is to derive the username from the assertion. If the field is left blank, it uses the string received in the NameID field of the incoming assertion as the username.</p> <p>If you choose a certificate attribute with more than one value, the system uses the first matched value. For example, if you enter <certDN.OU> and the user has two values for the attribute (ou=management, ou=sales), the system uses "management". To use all values, add the SEP attribute to the variable. For example, if you enter <certDN.OUT SEP=":">, the system uses "management:sales". The attributes received in the attribute statement in the incoming assertion are saved under userAttr. These variables can also be used with angle brackets and plain text. If the username cannot be generated using the specified template, the login fails. If the NameID field of the incoming assertion is of type X509Nameformat, then the individual fields can be extracted using system variable "assertionNameDN".</p> <p>Currently supported NameIDs are - EMAIL, X509_SUBJECT, WIN_DOMAIN_QUALIFIED. If a SAML request is received with a different NameID format, then processing of the request fails with unsupported NameID format error message.</p>
Allowed Clock Skew (minutes)	<p>Specify the maximum allowed difference in time between the system clock and the SAML identity provider server clock.</p> <p>SAML is a time sensitive protocol. The time-based validity of a SAML assertion is determined by the SAML identity provider. If the SAML identity provider and SAML service provider clocks are askew, the assertion can be determined invalid, and you will receive the following error:</p> <p>"SAML Transferred failed. Please contact your system administrator. Detail: Failure: No valid assertion found in SAML response."</p> <p>We recommend you use NTP to ensure the clocks are synchronized and that you set an Allowed Clock Skew value that accommodates any expected or permissible skew.</p>

Settings	Guidelines
Support Single Logout	<p>Single logout is a mechanism provided by SAML for logging out a particular user from all the sessions created by the identity provider. Select this option if the system must receive and send a single logout request for the peer SAML identity provider.</p> <p>If you use the metadata option, the Single Logout Service URL setting can be completed by selecting the SLO service URL from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page. The system sends Single Logout requests to this URL.</p> <p>In addition, if you use the metadata option, the Single Logout Response URL setting is completed based on your selection for Single Logout Service URL. If the identity provider has left this setting empty in its metadata file, the system sends the Single Logout response to the SLO service URL.</p> <p>If you complete these settings manually, ask the SAML identity provider administrator for guidance.</p> <p>The Support Single Logout service for the identity provider must present a valid certificate.</p>
SSO Method	

Settings	Guidelines
Artifact	<p>When configured to use the Artifact binding, the system contacts the Artifact Resolution Service (ARS) to fetch the assertion using SOAP protocol. If the ARS is hosted on a HTTPS URL, then the certificate presented by the ARS is verified by the system. For this verification to pass successfully, the CA of the server certificate issued to the identity provider ARS must be added to the trusted server CA on the system.</p> <p>Complete the following settings to configure SAML using the HTTP Artifact binding:</p> <p>Source ID. Enter the source ID for the identity provider ARS. Source ID is Base64-encoded, 20-byte identifier for the identity provider ARS. If left blank, this value is generated by the system.</p> <p>Source Artifact Resolution Service URL. For metadata-based configuration, this field is completed automatically from the metadata file and is not configurable. For manual configurations, enter the URL of the service to which the SP ACS is to send ArtifactResolve requests. ArtifactResolve requests are used to fetch the assertion from the artifact received by it.</p> <p>SOAP Client Authentication. Select HTTP Basic or SSL Client Certificate and complete the related settings. If you use an SSL client certificate, select a certificate from the device certificate list.</p> <p>Select Device Certificate for Signing. Select the device certificate the system uses to sign the AuthnRequest sent to the identity provider SSO service. If you do not select a certificate, the system does not sign AuthnRequest.</p> <p>Select Device Certificate for Encryption. Select the device certificate the system uses to decrypt encrypted data received in the SAML response. The public key associated with the device certificate is used by the identity provider for encryption.</p>

Settings	Guidelines
POST	<p>When configured to use the POST binding, the system uses a response signing certificate to verify the signature in the incoming response or assertion. The certificate file must be in PEM or DER format. The certificate you select should be the same certificate used by the identity provider to sign SAML responses.</p> <p>Complete the following settings to configure SAML using the HTTP POST binding:</p> <p>Response Signing Certificate. If you use the metadata-based configuration option, select a certificate from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page.</p> <p>If you configure these settings manually, browse to and upload the certificate to be used to validate the signature in the incoming response or assertion. If no certificate is specified, the certificate embedded in the response is used.</p> <p>Enable Signing Certificate status checking. Select this option to check the validity of the signing certificate before verifying the signature. This setting applies to any certificate used for signature verification. If this option is enabled, the response will be rejected if the certificate is revoked, expired, or untrusted. If this option is selected, the certificate CA must be added to the Trusted Client CA store.</p>
	<p>If this option is not enabled, then the certificate is used without any checks. Select Device Certificate for Signing. Select the device certificate the system uses to sign the AuthnRequest sent to the identity provider SSO service. If you do not select a certificate, the system does not sign AuthnRequest.</p>
	<p>Select Device Certificate for Encryption. Select the device certificate the system uses to decrypt encrypted data received in the SAML response. The public key associated with the device certificate is used by the identity provider for encryption.</p>

Settings	Guidelines
Authentication Context Classes	<p>Use the Add and Remove buttons to select authentication context classes to be sent in the authentication requests to the SAML identity provider. These are included in the RequestedAuthnContext element.</p> <p>In the OASIS standard, an authentication context is defined as "the information, additional to the authentication assertion itself, that the relying party may require before it makes an entitlements decision with respect to an authentication assertion."</p> <p>This feature supports all authentication context classes specified in the SAML 2.0 OASIS Authn Context specification.</p> <p>For example, if you select X509, the system sends the following context:</p> <pre><samlp:RequestedAuthnContext> <saml:AuthnContextClassRef xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion"> urn:oasis:names:tc:SAML:2.0:ac:classes:X509</saml:AuthnContextClassRef> </samlp:RequestedAuthnContext></pre> <p>In response, the SAML IdP sends the context data along with the authentication results. The system stores the context data in the session cache and as a system variable named <code>samlAuthnContextClass</code>. The system variable can be used in role mapping rules and resource policy detailed rules.</p> <p>Specify a comparison attribute within the RequestedAuthnContext element. The comparison attribute specifies the relative strengths of the authentication context classes specified in the request and the authentication methods offered by a SAML IdP. The following values defined in the SAML 2.0 OASIS core specification can be selected:</p> <ul style="list-style-type: none"> exact - Requires the resulting authentication context in the authentication statement to be the exact match of at least one of the authentication contexts specified. minimum - Requires the resulting authentication context in the authentication statement to be at least as strong as one of the authentication contexts specified. maximum - Requires the resulting authentication context in the authentication statement to be stronger than any one of the authentication contexts specified. better - Requires the resulting authentication context in the authentication statement to be as strong as possible without exceeding the strength of at least one of the authentication contexts specified. <p>Select the same value that is configured on the SAML IdP. If none is specified in the SAML IdP configuration, the implicit default is exact.</p>

Settings	Guidelines
Service Provider Metadata Settings	
Metadata Validity	Enter the number of days the metadata is valid. Valid values are 0 to 9999. 0 specifies the metadata does not expire.
Do Not Publish SA Metadata	Select this option if you do not want to publish the metadata at the location specified by the Entity ID field.
Download Metadata	This button appears only after you have saved the authentication server configuration. Use this button to download the metadata of the current SAML service provider.
User Record Synchronization	
Enable User Record Synchronization	Allow users to retain their bookmarks and individual preferences regardless of which device they log in to.
Logical Auth Server Name	Specify the server name if you have enabled user record synchronization.

Using a Time-Based One-Time Password (TOTP) Authentication Server

This topic describes the Ivanti Connect Secure's integration with the Time-Based One-Time Password (TOTP) Authentication Servers.

TOTP Authentication Server Overview

This section describes support for using the Local/Remote Ivanti Connect Secure TOTP authentication server.

Understanding TOTP

Time-based One-Time Password (TOTP) algorithm as defined in RFC6238 is an authentication mechanism where a one-time password (a.k.a token) is generated by the authentication server and client from a shared secret key and the current time. ICS can act as TOTP authentication server. Any third-party TOTP applications (for example, Windows Authenticator or Google Authenticator) available on the mobile and desktop client platforms generate TOTP tokens. The TOTP authentication option is natively available on ICS without any additional products or license requirements. Customers can use TOTP authentication as part of their MFA policy, and strengthen their authentication mechanism for secure access scenarios.

Interoperability Requirements and Limitations

Before you begin:

- TOTP authentication server users' configuration is automatically synchronized within all nodes in a single cluster. If there are multiple clusters behind a DNS load-balancer, then the admin has to manually perform binary export/import user's configuration to all the nodes in different clusters.
- TOTP feature is configurable across clusters.
- First time users have to register a new TOTP user-account via web. End-users cannot use Pulse Desktop applications and Pulse Mac applications for new user registration.

Pre-9.0R3 users with more than one TOTP account will get reset when the system software is upgraded to ICS 9.0R3 or later. In such case, users have to re-register with TOTP.

- Two standalone nodes or separate clusters can be synced. For now, binary import/export of user configuration option can be used.



For the users who are already using custom sign-in pages:


For TOTP authentication to work, existing custom sign-in pages need to include following sign-in pages:

- TotpAuthRegister.shtml
- TotpAuthRegister-mobile-webkit.shtml
- TotpAuthRegister-ipad.shtml
- TotpAuthRegister-stdaln.shtml
- TotpAuthTokenEntry.shtml

- TotpAuthTokenEntry-mobile-webkit.shtml
- TotpAuthTokenEntry-ipad.shtml
- TotpAuthTokenEntry-stdIn.shtml

These files can be downloaded from sample custom sign-in pages URL: <https://<ICS>/dana-admin/download/sample.zip?url=/dana-admin/auth/custompage.cgi?op=Download&samplePage=sample>

Configuring Authentication with a TOTP Authentication Server

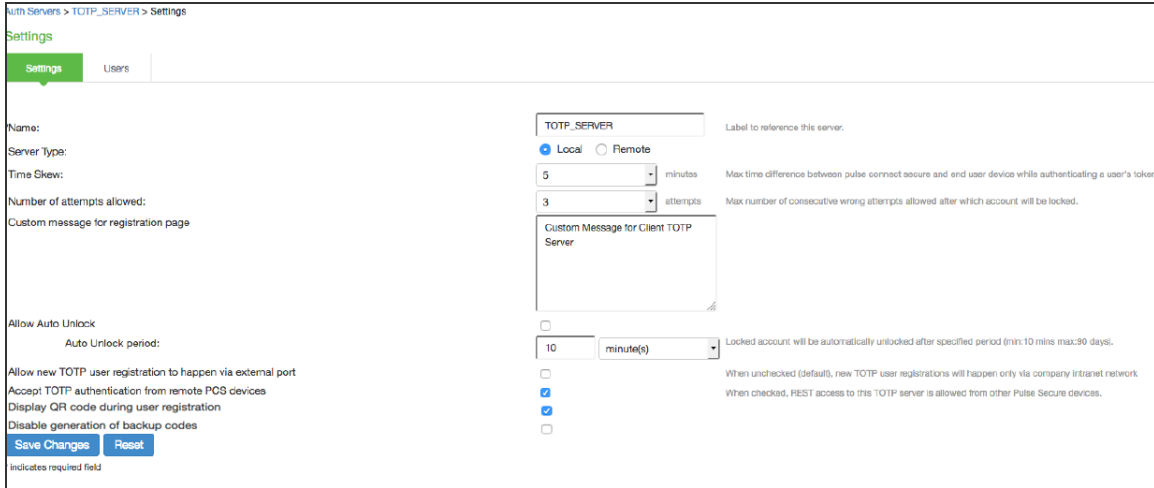


Configuring the TOTP Authentication Server Settings

To configure the TOTP server as Local:

1. Select **Authentication > Auth. Servers**.
2. Select **Time based One-Time Password (TOTP) Server** and click **New Server** to display the configuration page.
3. Complete the configuration as described in the following table.
4. Save the configuration.

The following figure depicts the TOTP Authentication Server Page - Local:



The following table lists the TOTP Auth Server Settings - Local

Settings	Guidelines
Name	Specify a name to identify the server within the system.
Server Type	TOTP server can be configured as local or remote. Select Local. Local: TOTP context is created locally and user database is maintained locally on the same device.
Time Skew	Specify maximum time difference between Ivanti Connect Secure and end user device while authenticating a user's token. (minimum: 1 minute, maximum: 5 minutes).
Number of attempts allowed	Specify maximum number of consecutive wrong attempts allowed after which account will be locked (minimum: 1 attempt, maximum: 5 attempts).
Custom message for registration page	Specify a custom message which can be shown on new TOTP user registration web-page.
Allow Auto Unlock	When checked, locked account will be automatically unlocked after specified period. (minimum: 10 minutes, maximum: 90 days)
Allow new TOTP user registration to happen via external port	When unchecked (default), new TOTP user registrations will happen only via internal port

Settings	Guidelines
Accept TOTP authentication from remote ICS devices	When checked, REST access to this TOTP server is allowed from other Ivanti Connect Secure devices.
Display QR code during user registration	When checked, displays QR code during user registration.
Disable generation of backup codes	When unchecked, generates backup codes.

To configure the TOTP server as Remote:

1. Select **Authentication > Auth. Servers**.
2. Select **Time based One-Time Password (TOTP) Server** and click **New Server** to display the configuration page. See Figure 16.
3. Complete the configuration as described in the following table.
4. Save the configuration.



If ICS is configured to use Remote TOTP server, then the remote ICS should have a valid certificate issued by a Trusted CA.

The following figure depicts the TOTP Authentication Server Page - Remote:

The following table lists TOTP Auth Server Settings - Remote:

Settings	Guidelines
Name	Specify a name to identify the server within the system.
Server Type	TOTP server can be configured as local or remote. Select Remote. Remote: In this configuration, authentication checks take place on the remote TOTP server. The user local device (ICS to which user is logging in) will act merely as a proxy between the user's client device and TOTP server. The communication to the remote device happens on REST API.
Allow new TOTP user registration to happen via external port	If this option is not selected, new TOTP user registrations happen only via company intranet network.
Host Name/IP	Specify remote host name or IP address where the TOTP server is configured. The IP address or host name must match the common name mentioned in the remote TOTP server certificate.
TOTP Server Name	This is the name of the TOTP server configured on the Remote TOTP server.
REST API Login	Enter the REST API login name.
REST API Password	Enter the REST API password.
REST Authentication Realm	Enter the realm name, which refers to the realm that should be used for authenticating the REST user (using the auth. server mapped to the Realm). WARNING: In 9.1R1 this field is mandatory. If the realm field is not entered, user logins fail after upgrade to 9.1R1.
Test Connection	This button is used to validate the connection to the remote TOTP server.



Customer needs to upload proper certificate to the Remote TOTP server. Wildcard certificate is also supported.

Configuring Admin/User Realm to Associate a TOTP Authentication Server as Secondary Authentication Server

For example, to configure a user realm:

1. Select **Users > User Realms > New User Realm**.
2. Complete the settings for the user-realm.

3. Check the **Enable additional authentication** server option.
4. Under Additional Authentication Server, select any already created **TOTP** authentication-server from the Authentication #2 dropdown, as shown in the following figure.



Whenever admin selects TOTP authentication-server as the additional authentication server, then the Username: Predefined as **<USER>** and Password: specified by user in sign-in page options are set by default.

5. Click on **Save Changes**.

The following figure depicts Configuring Admin/User Realm to Associate a TOTP Auth. Server as Secondary Auth. Server:

The screenshot shows the 'General' tab of the 'Users' realm configuration. The 'Additional Authentication Server' section is expanded, showing the following settings:

- Enable additional authentication server
- Authentication #2: Demo TOTP
- Username is:
 - specified by user on sign-in page
 - predefined as: <USER>
 - specified by user on sign-in page
 - predefined as: <PASSWORD>
- Password is:
 - specified by user on sign-in page
 - predefined as: <PASSWORD>
- End session if authentication against this server fails

Other visible settings include:

- Name: Users
- Description: Default authentication realm for users
- When editing, start on the Role Mapping page:
- Authentication: System Local
- User Directory/Attribute: None
- Accounting: None
- Device Attributes: None
- Dynamic policy evaluation: Enable dynamic policy evaluation

Using Google Authenticator Application to Register to a TOTP Server

The admin can associate an end-user to a realm that has a secondary authentication server configured as TOTP authentication server.

For first time registration via web, perform the following steps:

For example: Admin associates an end-user User1 to a user-realm that has the TOTP authentication-server configured as the secondary authentication-server.

When User1 for the first time, performs a log in to the above configured user-realm:

1. After successful authentication with primary authentication-server, User1 is shown the TOTP registration page. See the following figure.
2. User1 is given a TOTP registration key in text form/QR image form and 10 backup codes. User saves 10 backup codes in a safe place for using it later during authentication when end-user device (where Google Authenticator app is installed) is not available (in emergency).
3. Now, User1 opens the device where Google Authenticator app is installed, then either scans the QR image (or) manually adds a new user (for example: GA-User1) by entering the above given secret registration key.
4. The Google-Authentication app (for GA-User1) generates a new 6-digit number called as a token once in every 30 seconds.
5. Enter the current token in the registration page. Click on Sign In. On successful authentication with that token, User1 will be taken to his/her home page.

The following figure depicts First Time Registration to a TOTP Server:

Add user1 user account to your two factor authentication app

You will need to install a two factor authentication application on your smartphone or tablet.

1. Configure the App:

Open your two factor authentication app and add **user1** user account by scanning the below QR code.

If you can't use QR code, then enter [this text](#).



2. Store Backup Codes:

Backup codes can be used to access your account in the event you lose access to your device and cannot receive two-factor authentication codes. Following backup codes are only for one time use, we recommend you to save them securely.

IUQXKS	YG7ZTC
QFWEWV	VZK3GK
ZJ3L42	IN6ACJ
DNCSYG	60DWYT
ENO20P	MNYT2Z

3. Enter token code that the application generates:

For already registered user, perform the following steps:

1. The already-registered user (For example: User1), whose realm was associated with secondary authentication server configured as TOTP authentication server, accesses ICS URL via web (User1 has already registered TOTP user in Google Authenticator app.)
2. After successful authentication with primary authentication server, user1 is shown TOTP Token entry page as seen in Figure 19.
3. User1 opens Google Authentication app that was installed in mobile (or PC), enters the current token to the Authentication Code. If mobile is not available, user can enter any of the unused backup codes.
4. On successful authentication with the token, User1 can enter any of the unused backup codes.



A backup code can be used only once to successfully authenticate with the TOTP authentication server. Once used, the same backup code cannot be reused.

The following figure depicts the Google Authentication Token:

A screenshot of a web page titled "Two-Factor Authentication". The page has a light yellow background. At the top, there is a header with the title "Two-Factor Authentication". Below the header, there is a paragraph of text: "Open the two-factor authentication app on your device to view your authentication code and verify your identity . Currently if you do not have access to your device, use one of the backup codes saved previously." Below this text, there is a label "Authentication code:" followed by a text input field. At the bottom of the form, there is a "Sign In" button.

Displaying the User Accounts Table

To display user accounts:

1. Select **Authentication > Auth. Servers**.
2. Click the link for the authentication server you want to manage.
3. Click the Users tab to display the user accounts table. The user accounts table includes entries for the accounts that have been created. See the following figure.
 - The "Last Attempted" column shows the last time and date a user attempted to log in.
 - The "Last Successful Login" shows the last successful sign-in date and time for each user.
 - Under the "User Information" column, there are details available for a user's "Realm", "Primary AuthServer" and the "Status" columns

There are 3 possible states for the "Status" column:

- Active: TOTP user's account is in use (that is user has used this account less than stale period of this TOTP authentication server)
- Locked: TOTP user account has been locked due to maximum number of wrong login attempts
- Unregistered: TOTP user has seen registration page, but yet to complete the registration by entering the correct token in the registration page.

4. Use the controls to search for users and manage user accounts:

- To search for a specific user, enter a username in the Show users named field and click Update.

You can use an asterisk (*) as a wildcard, where * represents any number of zero or more characters. For example, to search for all usernames that contain the letters jo, enter *jo*. The search is case-sensitive. To display the entire list of accounts again, type * or delete the field's contents and click **Update**.

- To limit the number of users displayed on the page, enter a number in the Show N user's field and click **Update**.
- To unlock a user, select the specific user and click **Unlock**.
- To reset a user's credentials, select the specific user and click **Reset**.

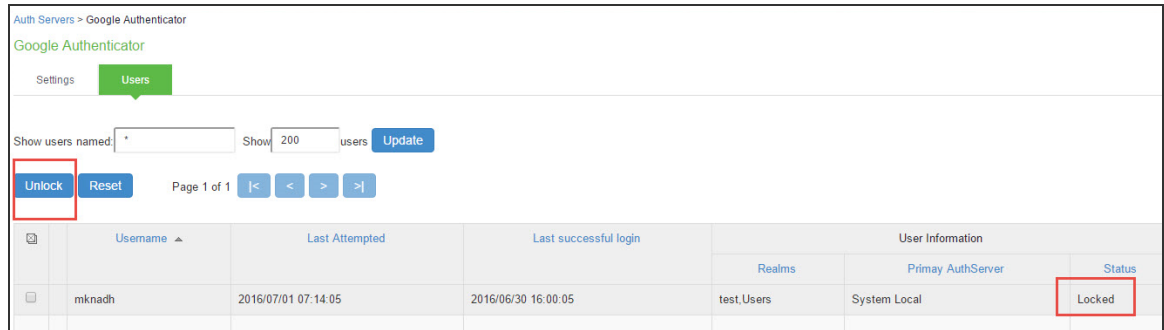
The following figure depicts Displaying the User Accounts Table:

	Username	Last Attempted	Last successful login	User Information		
				Realms	Primary AuthServer	Status
<input type="checkbox"/>	flower.com\user1	2016/06/28 11:55:30	2016/06/28 11:54:46	Users	AD	Active
<input type="checkbox"/>	flower.com\user4	2016/06/27 09:20:27	2016/06/24 15:55:08	Users	AD	Active
<input type="checkbox"/>	user1	2016/06/24 10:01:26	2016/06/24 09:25:59	test	System Local	Active
<input type="checkbox"/>	flower.com\user1	2016/06/24 10:02:14	2016/06/23 15:44:43	Users	AD	Locked
<input type="checkbox"/>	flower.com\user5	2016/06/24 09:09:15	2016/06/21 14:20:18	Users	AD	Active
<input type="checkbox"/>	user2	2016/06/24 10:01:33		test	System Local	Unregistered
<input type="checkbox"/>	user3	2016/06/24 10:01:43		test	System Local	Unregistered
<input type="checkbox"/>	flower.com\ycsunil	2016/06/28 12:37:09		Users	AD	Unregistered
<input type="checkbox"/>	flower.com\user3	2016/06/27 09:20:30		Users	AD	Unregistered

To unlock a TOTP user's account:

1. At **Authentication > Auth Servers > Google Authenticator** - Go to the Users tab. The list of users is displayed.
2. Select the user whose account you choose to unlock.
3. Click on the **Unlock** button.

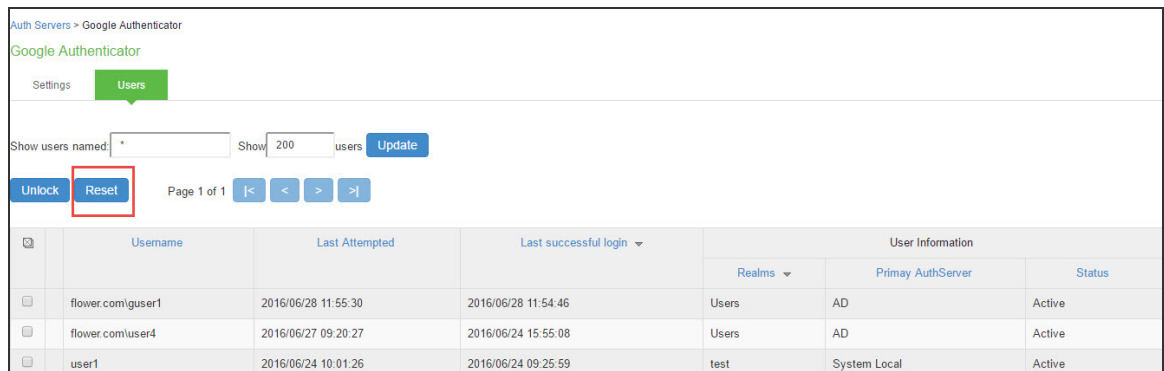
The following figure depicts Unlocking a User:



To reset a TOTP user's account:

1. At **Authentication > Auth Servers > Google Authenticator** - Go to the **Users** tab. The list of users is displayed.
2. Select the user whose account you choose to reset.
3. Click on the **Reset** button. This removes the user entry from the table.

The following figure depicts Resetting a User:

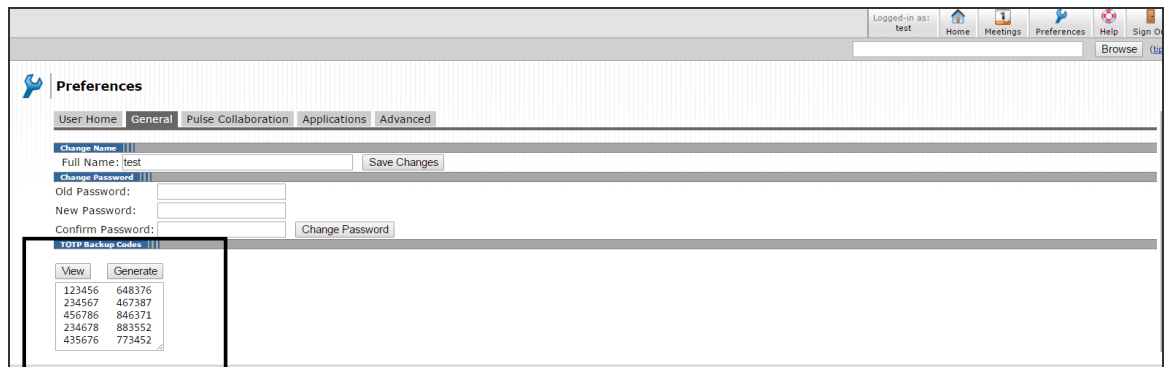
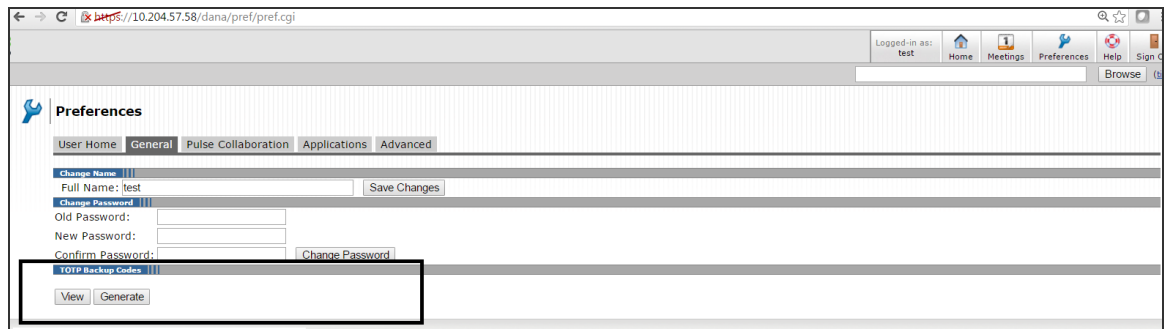


Viewing/Generating Backup Codes

To view/generate TOTP backup codes after successful log in to a TOTP server via web:

1. User successfully authenticates to primary auth-server and TOTP auth-server via web.
2. Click on the Preference option on the top of the page.
3. In the Preference page, under TOTP Backup codes, click on either View or Generate to obtain user's TOTP backup codes.

The following figure depicts View/Generate Backup Codes:



Exporting/Importing TOTP Users

To export/import TOTP users:

1. Select **Authentication > Auth. Servers**.
2. Click the link for the authentication server you want to manage.
3. Click the Users tab to display the user accounts table. The user accounts table includes entries for the accounts that have been created. See the following figure.
4. Use the Export and Import buttons located at the bottom of the user accounts table to export and import TOTP users data.

Authentication Realms

Understanding Authentication Realms

An authentication realm specifies the conditions that users must meet in order to sign into the system. A realm consists of a grouping of authentication resources, including:

- An authentication server - verifies that the user is who he claims to be. The system forwards credentials that a user submits on a sign-in page to an authentication server.
- A directory server - an LDAP server that provides user and group information to the system that the system uses to map users to one or more user roles.
- An authentication policy - specifies realm security requirements that need to be met before the system submits a user's credentials to an authentication server for verification.
- Role mapping rules - conditions a user must meet in order for the system to map the user to one or more user roles. These conditions are based on either user information returned by the realm's directory server or the user's username.

Authentication realms are an integral part of the access management framework, and therefore are available on all Ivanti Connect Secure products.

Creating an Authentication Realm

To create an authentication realm:

1. In the admin console, choose **Administrators > Admin Realms or Users > User Realms**.
2. On the respective **Authentication Realms** page, click **New**. Or, select a realm and click **Duplicate** to base your realm on an existing realm.
3. Enter a name to label this realm and (optionally) a description.
4. If you are copying an existing realm, click **Duplicate**. Then, if you want to modify any of its settings, click the realm's name to enter into edit mode.
5. Select **When editing, start on the Role Mapping page** if you want the Role Mapping tab to be selected when you open the realm for editing.
6. Under Servers, specify:
 - An authentication server to use for authenticating users who sign in to this realm.

- A directory/attribute server to use for retrieving user attribute and group information for role mapping rules and resource policies. (optional)
- A RADIUS accounting server to use to track when a user signs-in and signs-out of the Ivanti Connect Secure (optional).

7. If you want to submit secondary user credentials to an SSO-enabled resource or enable two-factor authentication to access the device, select Additional authentication server. Then:

1. Select the name of the secondary authentication server.
2. Select Username is specified by user on sign-in page if you want to prompt the user to manually submit his username to the secondary server during the sign-in process. Otherwise, if you want to automatically submit a username to the secondary server, enter static text or a valid variable in the predefined as field. By default, the system submits the <username> session variable, which holds the same username used to sign in to the primary authentication server.
3. Select Password is specified by user on sign-in page if you want to prompt the user to manually submit his password to the secondary server during the sign-in process. Otherwise, if you want to automatically submit a password to the secondary server, enter static text or a valid variable in the predefined as field.

Mask Static Password: From 8.3R4, a check box has been added to mask static password. This new check box by default is disabled and any new upgrade with this feature will show the UI as unchecked. Once password is masked, there is no way the password can be unmasked and only way would be to edit and set a new password. This option can be done for both Admin realm and User Realm.

4. Select End session if authentication against this server fails if you want to control access to the system based on the successful authentication of the user's secondary credentials. If selected, authentication fails if the user's secondary credentials fail.
5. If you want to use dynamic policy evaluation for this realm select Dynamic policy evaluation to enable an automatic timer for dynamic policy evaluation of this realm's authentication policy, role mapping rules, and role restrictions. Then:

1. Use the Refresh interval option to specify how often you want the Ivanti Connect Secure to perform an automatic policy evaluation of all currently signed-in realm users. Specify the number of minutes (5 to 1440).
 2. Select Refresh roles to also refresh the roles of all users in this realm. (This option does not control the scope of the Refresh Now button.)
 3. Select Refresh resource policies to also refresh the resource policies for all users in this realm. (This option does not control the scope of the Refresh Now button.)
 4. Click **Refresh Now** to manually evaluate the realm's authentication policy, role mapping rules, role restrictions, user roles, and resource policies of all currently signed-in realm users. Use this button if you make changes to an authentication policy, role mapping rules, role restrictions, or resource policies and you want to immediately refresh the roles of this realm's users.
8. Click **Save Changes** to create the realm on the device. The General, Authentication Policy, and Role Mapping tabs for the authentication realm appear.
 9. Perform the next configuration steps:
 1. Configure one or more role mapping rules.
 2. Configure an authentication policy for the realm.

Role Mapping Rules

Role mapping rules are conditions a user must meet in order for the system to map the user to one or more user roles. These conditions are based on either user information returned by the realm's directory server or the user's username. You must specify role mapping directives in the following format: If the specified condition is not true, then map the user to the selected roles.

You create a role mapping rule on Role Mapping tab of an authentication realm. When you click **New Rule** on this tab, the Role Mapping Rule page appears with an inline editor for defining the rule. This editor leads you through the three steps of creating a rule:

- Specify the type of condition on which to base the rule. Options include:

- Username
- User attribute
- Certificate or certificate attribute
- Group membership
- Custom expressions
- Specify the condition to evaluate, which consists of:
 - One or more usernames, user attributes, certificate attributes, groups (LDAP), or expressions depending on the type of condition you selected.
 - To what the value(s) should equate, which may include a list of usernames, user attribute values from a RADIUS or LDAP server, client-side certificate values (static or compared to LDAP attributes), LDAP groups, or predefined custom expressions.
- Specify the roles to assign to the authenticated user.

The system compiles a list of eligible roles to which a user may be mapped, which are roles specified by the role mapping rules to which the user conforms. Next, the system evaluates the definition for each role to determine if the user complies with any role restrictions. The system uses this information to compile a list of valid roles, which are roles for which the user meets any additional requirements. Finally, the system either performs a permissive merge of the valid roles or presents a list of valid roles to the user, depending on the configuration specified on the realm's Role Mapping tab.

Specifying Role Mapping Rules for an Authentication Realm

When creating a new rule that uses LDAP user attributes, LDAP group information, or custom expressions, you must use the server catalog.

To specify role mapping rules for an authentication realm:

1. In the admin console, choose **Administrators > Admin Realms or Users > User Realms**.
2. On the respective Authentication Realms page, select a realm and then click the **Role Mapping** tab.
3. Click **New Rule** to access the Role Mapping Rule page. This page provides an inline editor for defining the rule.
4. In the Rule based on list, choose one of the following:

- **Username** - Username is the system username entered on the sign-in page. Choose this option if you want to map users to roles based on their system usernames. This type of rule is available for all realms.
- **User attribute** - User attribute is a user attribute from a RADIUS, LDAP server. Choose this option if you want to map users to roles based on an attribute from the corresponding server. This type of rule is available only for realms that use a RADIUS server for the authentication server, or that use an LDAP server for either the authentication server or directory server. After choosing the User attribute option, click Update to display the Attribute list and the Attributes button. Click the Attributes button to display the server catalog.
- **Certificate or Certificate attribute** - Certificate or Certificate attribute is an attribute supported by the users' client-side certificate. Choose this option if you want to map users to roles based on certificate attributes. The Certificate option is available for all realms; the Certificate attribute option is available only for realms that use LDAP for the authentication or directory server. After choosing this option, click Update to display the Attribute text box.
- **Group membership** - Group membership is group information from an LDAP or native Active Directory server that you add to the server catalog Groups tab. Choose this option if you want to map users to roles based on either LDAP or Active Directory group information. This type of rule is available only for realms that use an LDAP server for either the authentication server or directory server or that use an Active Directory server for authentication. (Note that you cannot specify an Active Directory server as an authorization server for a realm.)
- **Custom Expressions** - Custom Expressions is one or more custom expressions that you define in the server catalog. Choose this option if you want to map users to roles based on custom expressions. This type of rule is available for all realms. After choosing this option, click Update to display the Expressions lists. Click the Expressions button to display the Expressions tab of the server catalog.



If you add more than one custom expression to the same rule, the system creates an "OR" rule for the expressions. For example, you might add the following expressions to a single rule:

- Expression 2: loginTime = (8:00AM TO 5:00PM)

Based on these expressions, a user would match this rule if he signed into the device between 8:00 and 5:00.

1. Under Rule, specify the condition to evaluate, which corresponds to the type of rule you select and consists of:

- Specifying one or more usernames, RADIUS or LDAP user attributes, certificate attributes, LDAP groups, or custom expressions.
 - Specifying to what the value(s) should equate, which may include a list of usernames, user attribute values from a RADIUS, or LDAP server, client-side certificate values (static or LDAP attribute values), LDAP groups, or custom expressions.
2. Then assign these roles:
 - Specify the roles to assign to the authenticated user by adding roles to the Selected Roles list.
 - Check Stop processing rules when this rule matches if you want to stop evaluating role mapping rules if the user meets the conditions specified for this rule.
 3. Click **Save Changes** to create the rule on the Role Mapping tab. When you are finished creating rules:

Make sure to order role mapping rules in the order in which you want to evaluate them. This task is particularly important when you want to stop processing role mapping rules upon a match.

Machine Authentication for Pulse Client Connections

Pulse client supports machine authentication. Machine authentication uses machine credentials (machine name and password or machine certificate) to authenticate the endpoint. You can enable machine authentication for the system as part of a Pulse Client Connection and distribute the connection to endpoints through the normal Ivanti distribution methods.

The following describes the requirements for a machine authentication environment:

- Machine authentication for Ivanti Connect Secure is available for Ivanti layer 3 connections only.
- The authentication server used by the Pulse client connection must be Active Directory/Windows NT for machine name/password authentication or a certificate server for machine certificate authentication. You can also use machine credentials when authenticating to RADIUS servers that verify the machine credentials against an Active Directory listing.
- The endpoint must be a member of a Windows domain and the machine credentials must be defined in Active Directory. Typically, during login, the user must enter domain/user in the username box.

- The Pulse client connection must be configured so that no prompts are presented during the login process. For example, prompts for realm or role selection or a server certificate trust prompt cause the connection to fail.
- For machine certificate authentication, the domain workstation logon certificate must be issued by the domain certificate authority. The root certificate (CA) must be in the Machine Trusted Certificate store instead of the certificate store for a particular user.

To enable a Ivanti connection for machine authentication:

1. Click **Users > Pulse clients > Connections** and create or select a connection set.
2. Create or edit a connection. Machine authentication is available for connection type Ivanti Connect Secure or Ivanti Policy Secure (L3), Ivanti Policy Secure (802.1X), or SRX only.
3. Under Connection is established, select one of the following options:
 - Automatically when the machine starts. Machine credentials used for authentication-This option enables machine-only authentication. Machine credentials are used to connect to the system before the user logs on. The user does not need to be logged in. The connection is maintained when a user logs on, logs off, or switches to a different logon.
 - Automatically when the machine starts. Connection is authenticated again when the user signs in into the desktop-This option enables user-after-desktop authentication. Machine credentials are used to authenticate the endpoint when no user is logged on. When a user logs on, the machine authentication connection is dropped, and the user login is used instead. When the user logs off, the machine connection is reestablished.

Connection Realm and Role Preferences for Machine Authentication

When a Pulse Client Connection is configured to use machine authentication, any prompts that occur during the login process cause the connection to fail. For example, if the Ivanti server authentication policy allows a user to select a realm or a role during the login process, Ivanti presents a dialog box to the user and prompts for the realm or role selection. To avoid failed connections due to prompts during machine authentication you can specify a preferred role and realm for a Ivanti connection.

For a Ivanti connection that is used for machine authentication, you do not need to specify the preferred role if any of the following conditions are true:

- Users are mapped to only one role.

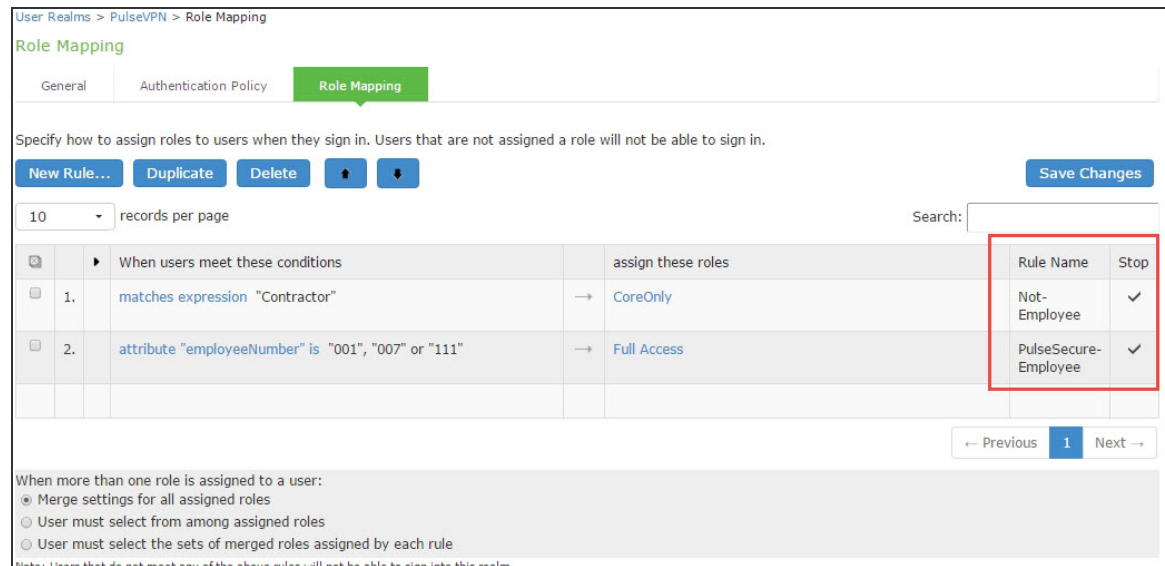
- Users are mapped to more than one role, but the realm's role mapping properties are set to merge settings for all assigned roles.

For a Ivanti connection that is used for machine authentication, you must specify the preferred realm if the authentication sign-in policy allows the user to select a realm. If that realm maps to only one role, you do not need to specify the role.

For a connection that is used for machine authentication, you must specify the preferred role if any of the following conditions are true:

- The realm that the user connects to maps to more than one role and the realm's role mapping properties are set to require that the user must select a role. The preferred role set must be the name of a role assigned in that realm.
- The realm that the user connects to maps to more than one role and the realm's role mapping properties are defined by role mapping rules. You specify the preferred role by specifying the name of a rule that assigns the role set. The following figure shows a role mapping rule with the rule name highlighted.

The following figure depicts the Pulse Client Role Mapping Rule:



When you create a connection for machine authentication, you must use the connection type Ivanti Connect Secure or Ivanti Policy Secure (L3), Ivanti Policy Secure (802.1X), or SRX. To identify the connection as a machine authentication connection, you specify how the connection is established using one of the following options:

- Automatically when the machine starts. Machine credentials used for authentication

This option uses the machine credentials defined in Active Directory for the machine login process and uses the same credentials for user login. When you select this option, the Realm and Role Set Preferences enable you to specify the following options:

- **Preferred Machine Realm**-Type the realm name that maps to the role you want to assign.
- **Preferred Machine Role Set**-Type the name of the role. The role must be one that is identified in the realm's role mapping properties. Or specify the name of a role mapping rule that assigns the role set.
- Automatically when the machine starts. Connection is authenticated again when the user signs in into the desktop

This option uses the Active Directory machine credentials for the machine login process. When machine login is complete, Ivanti drops that connection and then uses the user credentials for user login. When you select this option, the Realm and Role Set Preferences enable you to specify the following options:

- **Preferred Machine Realm**-Type the realm name that maps to the role you want to assign.
- **Preferred Machine Role Set**-Type the name of the role. The role must be one that is identified in the realm's role mapping properties. Or specify the name of a role mapping rule that assigns the role set.
- **Preferred User Realm**-Type the realm name that maps to the role you want to assign.
- **Preferred User Role Set**-Type the name of the role. The role must be one that is identified in the realm's role mapping properties. Or specify the name of a role mapping rule that assigns the role set.



Realm and role prompts are not the only prompts that are possible during the login process. If the connection has the Dynamic Certificate Trust option enabled, and there is an issue with the server certificate, a question asks the user if it is Ok to proceed. That certificate prompt causes a machine connection to fail. Note that the prompt for upgrading software is presented after the user connection is established and it will not affect a machine authentication connection.

Configuring Role Mapping Rules based on Geo Location Custom Expressions

An admin can configure role mapping rules for any realm based on Geo Location custom expressions.

To create a role mapping rule:

1. Select **Users > User Realms**.
2. On the User Realms page, select a realm and then click the **Role Mapping** tab.
3. Click **New Rule** to access the Role Mapping Rule page. This page provides an in-line editor for defining the rule.
4. From the Rule based on list, select '**Custom Expressions**' and click **Update**.
5. Enter an appropriate Name for the expression.
6. Click the **Expressions** tab. In the **Expressions Dictionary** box, under Variables look for 'geoLocationCountry'

The following figure depicts the geoLocationCountry Expression:

The screenshot shows the 'Expressions' tab in a web interface. At the top, there are tabs for 'Attributes', 'Groups', 'Expressions' (which is highlighted in green), and 'Variables'. Below the tabs, there is a 'View:' dropdown menu set to 'Allowed_Countries'. The 'Name:' field contains 'Allowed_Countries'. The 'Expression:' text area contains the code: `geoLocationCountry = ('United States' or 'Sri Lanka')`. To the right, the 'Expressions Dictionary' panel is visible, showing a list of variables. The 'geoLocationCountry' variable is selected, and its description is: 'Geo Location. The location from where user should be allowed or denied to login from.' Below the description, there is an 'Example:' section with the code: `geoLocationCountry = 'United States'` and `geoLocationCountry = ('United States' or 'Canada')`. Other variables listed include 'eapProtocol', 'group.<group-name>', 'groups', 'hostCheckerPolicy', and 'locationGrp'. At the bottom of the interface, there are three buttons: 'Save Changes', 'Close', and 'Delete'.


7. Copy the Examples text to the **Expression** box and change the **Country Name** of your choice.
8. Click **Add Expression** and then click **Close**.
9. Select the rule you just created from the **Available Expressions** list and click **Add** to move it to the **Selected Expressions** list.

Specify the roles to assign by adding roles to the **Selected Roles** list.

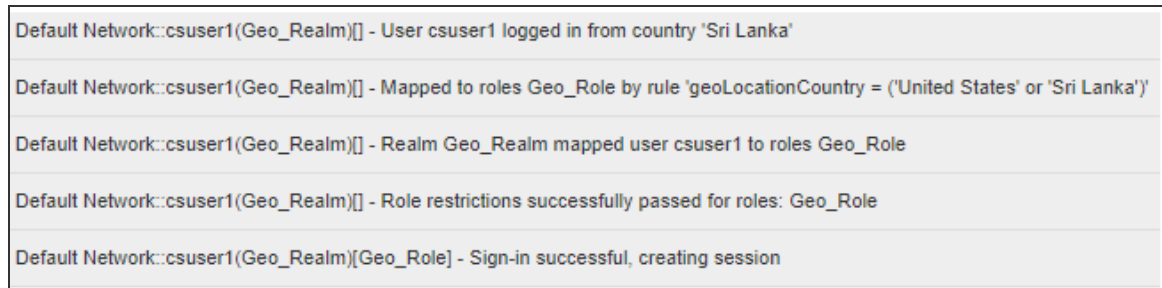
10. Click **Save Changes**.

The new rule will be listed in the Role Mapping list.

The following figure depicts geoLocationCountry Role Mapping Rule:



The following screen shows the Policy Trace log output where the role is mapped based on the defined Rules.



The list of countries supported:

Afghanistan	British Indian Ocean Territory	Egypt	Heard Island and McDonald Islands
Aland Islands	Brunei Darussalam	El Salvador	Holy See (Vatican City State)
Albania	Bulgaria	Equatorial Guinea	Honduras
Algeria	Burkina Faso	Eritrea	Hong Kong
American Samoa	Burundi	Estonia	Hungary
Andorra	Cambodia	Ethiopia	Iceland
Angola	Cameroon	Europe	India
Anguilla	Canada	Falkland Islands (Malvinas)	Indonesia
Antarctica	Cape Verde	Faroe Islands	Iran
Antigua and Barbuda	Cayman Islands	Fiji	Iraq
Argentina	Central African Republic	Finland	Ireland
Armenia	Chad	France	Isle of Man
Aruba		French Guiana	
Asia/Pacific Region			

Australia	Chile	French Polynesia	Israel
Austria	China	French Southern Territories	Italy
Azerbaijan	Christmas Island	Gabon	Jamaica
Bahamas	Cocos (Keeling) Islands	Gambia	Japan
Bahrain	Colombia	Georgia	Jersey
Bangladesh	Comoros	Germany	Jordan
Barbados	Congo	Ghana	Kazakhstan
Belarus	Cook Islands	Gibraltar	Kenya
Belgium	Costa Rica	Greece	Kiribati
Belize	Cote d'Ivoire	Greenland	Korea
Benin	Croatia	Grenada	Kuwait
Bermuda	Cuba	Guadeloupe	Kyrgyzstan
Bhutan	Curacao	Guam	Lao People's Democratic Republic
Bolivia	Cyprus	Guatemala	Latvia
Bonaire	Czech Republic	Guernsey	Lebanon
Bosnia and Herzegovina	Denmark	Guinea	Lesotho
Botswana	Djibouti	Guinea-Bissau	Liberia
Bouvet Island	Dominica	Guyana	Libyan Arab Jamahiriya
Brazil	Dominican Republic	Haiti	Liechtenstein
	Ecuador		
Lithuania	Nicaragua	Saint Vincent and the Grenadines	Tanzania
Luxembourg	Niger	Samoa	Thailand
Macao	Nigeria	San Marino	Timor-Leste
Macedonia	Niue	Sao Tome and Principe	Togo
Madagascar	Norfolk Island	Saudi Arabia	Tokelau
Malawi	Northern Mariana Islands	Senegal	Tonga
Malaysia	Norway	Serbia	Trinidad and Tobago
Maldives	Oman	Seychelles	Tunisia
Mali	Pakistan	Sierra Leone	Turkey
Malta	Palau	Singapore	Turkmenistan
Marshall Islands	Palestinian Territory	Sint Maarten	Turks and Caicos Islands
Martinique	Panama	Slovakia	Tuvalu
Mauritania	Papua New Guinea	Slovenia	Uganda
Mauritius	Paraguay	Solomon Islands	Ukraine
Mayotte	Peru	Somalia	United Arab Emirates
Mexico	Philippines	South Africa	United Kingdom
Micronesia	Pitcairn		United States
Moldova	Poland		United States Minor Outlying Islands
Monaco	Portugal		
Mongolia			

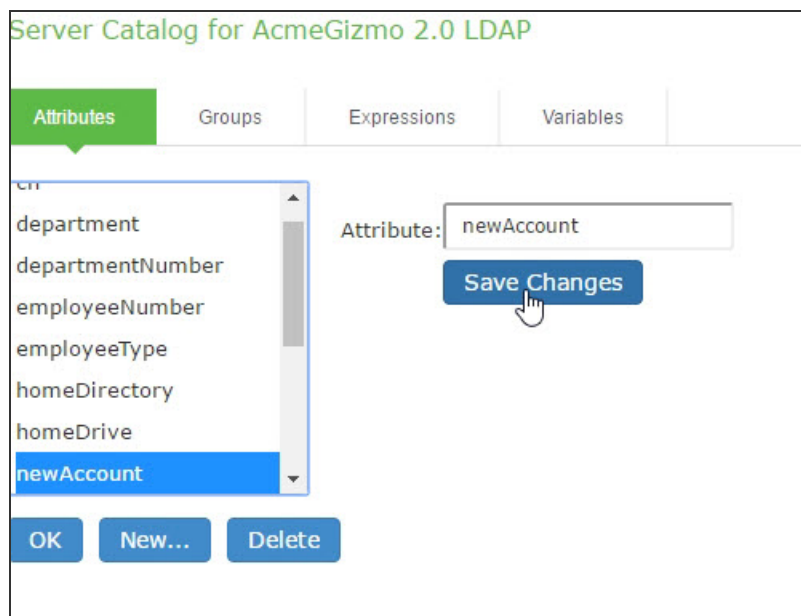
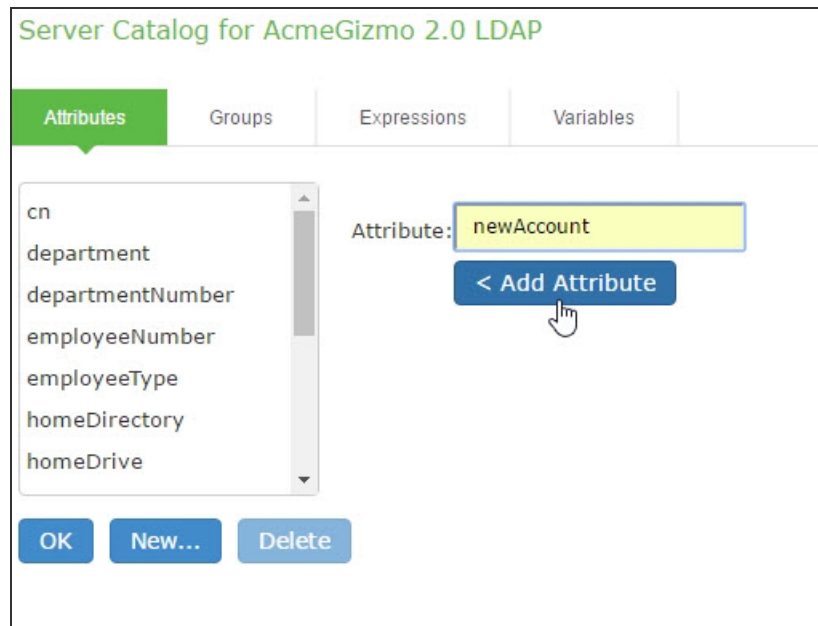
Montenegro	Puerto Rico	South Georgia and the South Sandwich Islands	Uruguay
Montserrat	Qatar	South Sudan	Uzbekistan
Morocco	Reunion	Spain	Vanuatu
Mozambique	Romania	Sri Lanka	Venezuela
Myanmar	Russian Federation	Sudan	Vietnam
Namibia	Rwanda	Suriname	Virgin Islands
Nauru	Saint Barthelemy	Svalbard and Jan Mayen	Wallis and Futuna
Nepal	Saint Helena	Swaziland	Western Sahara
Netherlands	Saint Kitts and Nevis	Sweden	Yemen
New Caledonia	Saint Lucia	Switzerland	Zambia
New Zealand	Saint Martin	Syrian Arab Republic	Zimbabwe
	Saint Pierre and Miquelon	Taiwan	
		Tajikistan	

Using the LDAP Server Catalog

The LDAP server catalog is a secondary window through which you specify additional LDAP information for the system to use when mapping users to roles, including:

- **Attributes** - The Server Catalog Attributes tab shows a list of common LDAP attributes, such as cn, uid, uniquemember, and memberof. This tab is accessible only when accessing the Server Catalog of an LDAP server. You can use this tab to manage an LDAP server's attributes by adding custom values to and deleting values from its server catalog. Note that the system maintains a local copy of the LDAP server's values; attributes are not added to or deleted from your LDAP server's dictionary. The following figure shows an example of adding the newAccount attribute.

The following figure depicts the Server Catalog > Attributes Tab - Adding an Attribute for LDAP:



- **Groups** - The Server Catalog Groups tab provides a mechanism to easily retrieve group information from an LDAP server and add it to the server catalog. You specify the BaseDN of your groups and optionally a filter to begin the search. If you do not know the exact container of your groups, you can specify the domain root as the BaseDN, such as dc=test, dc=test. The search page returns a list of groups from your server, from which you can choose groups to enter into the Groups list.



The BaseDN value specified in the LDAP server's configuration page under [chapter](#) is the default BaseDN value. The Filter value defaults to (cn=*).

You can also use the Groups tab to specify groups. You must specify the Fully Qualified Distinguished Name (FQDN) of a group, such as cn=GoodManagers, ou=HQ, ou=test, o=com, c=US, but you can assign a label for this group that appears in the Groups list. Note that this tab is accessible only when accessing the Server Catalog of an LDAP server. [The following figure depicts Server Catalog > Groups Tab - Adding LDAP Groups](#): and the following figures showing examples of adding LDAP and Active Directory groups.

The following figure depicts Server Catalog > Groups Tab - Adding LDAP Groups:

Server Catalog for AcmeGizmo 2.0 LDAP

Attributes **Groups** Expressions Variables

To add a group, type a name and click Add Group. To edit an existing group, select it, make your Save Changes. When you are done, click OK.

AcmeGizmo Admins
Connect SA Admins
Contractors
Employees
Help Desk
Helpdesk
Partners

Name:
DN:
Type:

< Add Group

OK New... Delete Search...

Group search for AcmeGizmo 2.0 LDAP

To search the LDAP server, specify a base DN and a filter, and click Search.

Base DN:

Filter:

10 records per page

<input type="checkbox"/>	Matching DNs	Type
<input type="checkbox"/>	CN=Administrators,CN=Builtin,DC=acmegizmo,DC=com	static
<input checked="" type="checkbox"/>	CN=Users,CN=Builtin,DC=acmegizmo,DC=com	static
<input type="checkbox"/>	CN=Guests,CN=Builtin,DC=acmegizmo,DC=com	static

Server Catalog for AcmeGizmo 2.0 LDAP

Attributes **Groups** Expressions Variables

To add a group, type a name and click Add Group. To edit an existing group, select it, make your changes, and click Save Changes. When you are done, click OK.

Employees

- Help Desk
- Helpdesk
- Partners
- Pulse Architects
- rDirectory Admins
- rDirectory Employee Admins
- Users

Name:

DN: CN=Users,CN=Builtin,DC=acmegizmo,DC=com

Type: static

Server Catalog for Acmegizmo - ADNT

Attributes **Groups** Expressions Variables

To add a group, type a name and click Add Group. To edit an existing group, select it, make your changes and click Save Changes. When you are done, click OK.

(none)

Name:

Group:

Enter group as DOMAIN/GroupName

< Add Group

OK New... Delete Search...

The following figure depicts the Server Catalog > Groups Tab - Adding Active Directory Groups:

Group search for Acmegizmo - ADNT

Add Selected Back

10 records per page Search:

<input type="checkbox"/>	Domain Groups
<input type="checkbox"/>	ACMEGIZMO/Enterprise Read-only Domain Controllers
<input type="checkbox"/>	ACMEGIZMO/Domain Admins
<input checked="" type="checkbox"/>	ACMEGIZMO/Domain Users
<input type="checkbox"/>	ACMEGIZMO/Domain Guests
<input type="checkbox"/>	ACMEGIZMO/Domain Computers
<input type="checkbox"/>	ACMEGIZMO/Domain Controllers
<input type="checkbox"/>	ACMEGIZMO/Schema Admins
<input type="checkbox"/>	ACMEGIZMO/Enterprise Admins

- **Expressions** - The Server Catalog Expressions tab provides a mechanism to write custom expressions for the role mapping rule.

To display the LDAP server catalog:

- After choosing the User attribute option on the Role Mapping Rule page, click Update to display the Attribute list and the Attributes button.
- Click the Attributes button to display the LDAP server catalog. (You can also click Groups after choosing the Group membership option, or click Expressions after choosing the Custom Expressions option.)

Customizing User Realm UI Views

You can use customization options on the User Authentication Realms page to quickly view the settings that are associated with a specific realm or set of realms. For instance, you can view the role-mapping rules that you have associated with all your user realms. Additionally, you can use these customized views to easily link to the authentication policies, servers, role-mapping rules, and roles associated with a user realm.

To view a sub-set of data on the User Authentication Realms page:

1. Select one of the following options from the View menu:

- **Overview** - Displays the authentication servers and dynamic policy evaluation settings that you have set for the specified user realms. You may also use this setting to link to the specified server configuration pages.
 - **Authentication Policy** - Displays Host Checker restrictions that you have enabled for the specified user realms. You may also use this setting to link to the specified Host Checker configuration pages.
 - **Role Mapping** - Displays rule conditions and corresponding role assignments that you have enabled for the specified user realms. You may also use this setting to link to the specified rule conditions and role assignments configuration pages.
 - **Servers** - Displays authentication server names and corresponding types that you have enabled for the specified user realms. You may also use this setting to link to the specified server configuration pages.
 - **Roles** - Displays role assignments and corresponding permissive merge settings that you have enabled for the specified user realms.
2. Select one of the following options from the for list:
 - **All realms** - Displays the selected settings for all user realms.
 - **Selected realms** - Displays the selected settings for the user realms you choose. If you select this option, select one or more of the check boxes in the Authentication Realm list.
 3. Click **Update**.

Single Sign-On

About Single Sign-On

Single sign-on (SSO) is a process that allows pre-authenticated Ivanti Connect Secure users to access other applications or resources that are protected by another access management system without having to re-enter their credentials.

The system provides several integration mechanisms that allow you to configure SSO connections from the system to other servers, applications, and resources. SSO mechanisms include:

- Remote SSO-The system provides loose integration with any application that uses a static POST action within an HTML form to sign in users. You can configure the system to post system credentials, LDAP attributes, and certificate attributes to a Web-enabled application, as well as set cookies and headers, allowing users to access the application without re-authenticating.
- SAML-The system provides loose integration with selected access management systems that use the Security Assertion Markup Language (SAML) to communicate with other systems. You can enable users to sign in to the system and then sign in to and access resources protected by the access management system without re-authenticating. You can also enable users to sign in to another access management system and then access resources protected by the system, without re-authenticating.
- Basic authentication and NTLM intermediation to Intranet sites-The system allows you to automatically submit user credentials to other web sites and proxies within the same Intranet zone. When you enable basic authentication intermediation through the Users > Resource Profiles > Web Applications/Pages page of the admin console, the system submits the cached credentials to Intranet web sites whose hostnames end in the DNS suffix configured in the System > Network > Overview page. To maximize security, you may also configure the system to use base-64 encoding to protect the cached credentials.
- Active Directory server-The system allows you to automatically submit Active Directory SSO credentials to other web sites and Windows file shares within the same Intranet zone that are protected by native NTLM authentication. When you enable this option, the system submits cached credentials to NTLM-protected web sites whose hostnames end in the DNS suffix configured in the System > Network > Overview page of the admin console.
- Terminal Sessions-When you enable the Terminal Services feature for a role, you allow users to connect to applications that are running on a Windows terminal server or Citrix MetaFrame server without re-authenticating.

The system determines which credentials to submit to the SSO-enabled server, application, or resource based on the mechanism you use to connect. Most mechanisms allow you to collect user credentials for up to two authentication servers in the system sign-in page and then submit those credentials during SSO.

The remaining mechanisms use unique methods for enabling SSO from Ivanti Connect Secure to the supported application.

About Multiple Sign-In Credentials

When configuring an authentication realm, you can enable up to two authentication servers for the realm. Enabling two authentication servers allows you to require two different sets of credentials—one for Ivanti Connect Secure and another for your SSO-enabled resource—without requiring the user to enter the second set of credentials when accessing the resource. It also allows you to require two-factor authentication in order to access the device.

Task Summary: Configuring Multiple Authentication Servers

To enable multiple authentication servers:

1. Create authentication server instances through the Authentication > Auth. Servers page of the admin console.
2. Associate the authentication servers with a realm using settings in the following pages of the admin console:
 - Users > User Realms > *Select Realm* > General
 - Administrators > Admin Realms > *Select Realm* > General
3. (Optional) Specify password length restrictions for the secondary authentication server using settings in the following pages of the admin console:
 - Users > User Realms > *Select Realm* > Authentication Policy > Password
 - Administrators > Admin Realms > *Select Realm* > Authentication Policy > Password

Task Summary: Enabling SSO to Resources Protected by Basic Authentication

To enable single sign-on to Web servers and Web proxies that are protected by basic authentication, you must:

1. Specify a hostname that ends with the same prefix as your protected resource using settings in the **System > Network > Overview** page of the admin console. (The system checks the hostnames to ensure that it is only enabling SSO to sites within the same Intranet.)
2. Enable users to access Web resources, specify the sites to which you want the system to submit credentials, create autopolicies that enable basic authentication intermediation single sign-on, and create bookmarks to the selected resources using settings in the **Users > Resource Profiles > Web Application/Pages > [Profile]** page of the admin console.
3. If you want users to access Web servers through a proxy, configure the system to recognize the appropriate servers and proxies using settings in the following pages of the admin console:
 - Use settings in **Users > Resource Policies > Web > Web proxy > Servers** page to specify which Web servers you want to protect with the proxy.
 - Use settings in the **Users > Resource Policies > Web > Web proxy > Policies** page to specify which proxies you want to use and which servers (above) you want the proxies to protect. You may specify individual resources on the server or the entire server.

Enabling SSO to Resources Protected by NTLM



The system supports web proxies that perform NTLM authentication. However, the following case is not supported: a proxy exists between the system and the back-end server and the back-end server performs the NTLM authentication.

To enable single sign-on to Web servers, Windows file servers, and Web proxies that are protected by NTLM, you must:


1. Specify a hostname that ends with the same suffix as your protected resource using settings in the **System > Network > Overview** page of the admin console. (The system checks the hostnames to ensure that it is only enabling SSO to sites within the same Intranet.)
2. Enable users to access the appropriate type of resource (Web or file), specify the sites or servers to which you want the system to submit credentials, create autopolicies that enable NTLM single sign-on, and create bookmarks to the selected resources using settings in the following pages of the admin console:

- **Users > Resource Profiles > Web Application/Pages > [Profile]**
 - **Users > Resource Profiles > File Browsing Resource Profiles > [Profile]**
3. If you want users to access Web servers through a proxy, configure the system to recognize the appropriate servers and proxies using settings in the following pages of the admin console:
 4. Use settings in Users > Resource Policies > Web > Web proxy > Servers page to specify which Web servers you want to protect with the proxy.
 5. Use settings in the Users > Resource Policies > Web > Web proxy > Policies page to specify which proxies you want to use and which servers (above) you want the proxies to protect. You may specify individual resources on the server or the entire server.

Multiple Sign-In Credentials Execution

The following diagram illustrates the process that the system uses to collect and authenticate multiple user credentials and submit them to SSO-enabled resources. Each of the steps in the diagram are described in further detail in the sections that follow.

The following figure depicts Collecting and Submitting Credentials from Multiple Servers:



Step 1: Collect the User's Primary Credentials

When the user signs in to a device, the system prompts him to enter his primary server credentials. The system saves these credentials to submit to the SSO resource later, if necessary. Note that the system saves the credentials exactly as the user enters them-it does not pre-pend or append them with additional information such as the user's domain.

Step 2: Collect or Generate the User's Secondary Credentials

You may configure the system to either manually collect or automatically generate the user's secondary set of credentials. If you configure the system to:

- Manually collect the user's secondary credentials-The user must enter his secondary credentials directly after entering his primary credentials.
- Automatically generate the user's credentials-The system submits the values you specified in the administration console during setup. By default, the system uses the <username> and <password> variables, which hold the username and password entered by the user for the primary authentication server.

For example, you may configure an LDAP server as your primary authentication server and an Active Directory server as your secondary authentication server. Then, you may configure the system to infer the user's Active Directory username but require the user to manually enter his Active Directory password. When the system infers the Active Directory username, it simply takes the name entered for the LDAP server (for example, JDoe@LDAPServer) and resubmits it to the Active Directory (for example, JDoe@ActiveDirectoryServer).

Step 3: Authenticate the Primary Credentials

After the system collects all required credentials, it authenticates the user's first set of credentials against the primary authentication server. Then:

- If the credentials successfully authenticate, the system stores them in the <username> and <password> session variables and continues on to authenticate the secondary credentials.



If you authenticate against a RADIUS server that accepts dynamic, time-sensitive passwords, you may choose to not store user passwords using the system session variable.

- If the credentials do not successfully authenticate, the system denies the user access to the device.

Step 4: Authenticate the Secondary Credentials

After authenticating the primary credentials, the system authenticates the secondary credentials. Then:

- If the credentials successfully authenticate, the system stores them in the <username[2]> and <password[2]> session variables and allows the user access to the device. You may also access these variables using the syntax <username@SecondaryServer> and <password@SecondaryServer>.



If you authenticate against a RADIUS server that accepts dynamic, time-sensitive passwords, you may choose to not store user passwords using the system session variable.

- If the credentials do not successfully authenticate, the system does not save them. Depending on how you configure your authentication realm, the system may allow or deny the user access to a device if his secondary credentials do not successfully authenticate.

Step 5: Submit Credentials to an SSO-Enabled Resource

After the user successfully signs in to a device, he may try to access an SSO-enabled resource using a pre-configured bookmark or other access mechanism. Then, depending on which type of resource the user is trying to access, the system submits different credentials. If the user is trying to access a:

- Web SSO or Terminal Services resource-The system submits the credentials that you specify through the admin console, such as <username> (which submits the user's primary credentials to the resource) or <username[2]> (which submits the user's secondary credentials to the resource). Or, if the user has entered a different username and password through the end user console, the system submits the user-specified credentials.



The system does not support submitting ACE server, or certificate server credentials to a Web SSO or terminal services resource. If you configure the system to submit credentials from one of these types of primary authentication servers, it submits credentials from the user's secondary authentication server instead. If these credentials fail, the system prompts the user to manually enter his username and password.

- **Resource protected by a Web server, Windows server, or Web proxy that is using NTLM authentication-The system submits credentials to the backend server or proxy that is protecting the Web or file resource. Note that you cannot disable NTLM authentication through the system-If a user tries to access a resource that is protected by NTLM, the system automatically intermediates the authentication challenge and submits credentials in the following order:**
 1. (Windows file resources only) Administrator-specified credentials-If you create a resource profile that specifies credentials for a Windows file resource and the user then accesses the specified resource, the system submits the specified credentials.
 2. Cached credentials-If the system does not submit administrator-specified credentials or the credentials fail, it determines whether it has stored credentials for the specified user and resource in its cache. (See below for information about when the system caches credentials.) If available, the system submits its stored credentials.
 3. Primary credentials-If the system does not submit cached credentials or the credentials fail, it submits the user's primary system credentials provided that following conditions are true:

- The resource is in the same Intranet zone as the device (that is, the resource's hostname ends in the DNS suffix configured in the System > Network > Overview page of the admin console).
 - (Web proxies only) You have configured the system to recognize the Web proxy through settings in the Users > Resource Policies > Web > Web Proxy pages of the admin console.
 - The credentials are not ACE credentials.
 - (RADIUS credentials only) You specify in the RADIUS configuration page that the RADIUS server does not accept one-time passwords.
4. Secondary credentials-If the primary credentials fail, the system determines whether it has secondary credentials for the user. If available, the system submits the user's secondary credentials provided that the conditions described for primary credentials are true.
 5. Last-entered credentials-If the system does not submit secondary credentials or if the credentials fail, it determines whether it has stored credentials for the specified user and a different resource in its cache. (See below for information about when the system caches credentials.) If available, the system submits its stored credentials provided the conditions described for primary credentials are true.
 6. User-specified credentials (prompt)-If the system does not submit last-entered credentials or if the credentials fail, it prompts the user to manually enter his credentials in the intermediate sign-in page. If the user selects the "Remember password?" check box, the system caches the user-specified credentials and, if necessary, resubmits them when the user tries to access the same resource again. Note that when the system caches these credentials, it remembers the specific user and resource, even after the user signs out of the device.
- **Resource protected by a Web server or Web proxy using basic authentication-The system submits credentials in the following order to the backend server or proxy that is protecting the Web resource:**
 1. Cached credentials-If the system does not submit administrator-specified credentials or the credentials fail, it determines whether it has stored credentials for the specified user and resource in its cache. If available, the system submits its stored credentials.
 2. Primary credentials-If the system does not submit cached credentials or the credentials fail, it submits the user's primary system credentials provided that following conditions are true:

- The resource is in the same Intranet zone as the device (that is, the resource's hostname ends in the DNS suffix configured in the System > Network > Overview page of the admin console).
 - (Web proxies only) You have configured the system to recognize the Web proxy through settings in the Users > Resource Policies > Web > Web Proxy pages of the admin console.
 - The credentials are not ACE credentials.
 - (RADIUS credentials only) You specify in the RADIUS configuration page that the RADIUS server does not accept one-time passwords.
3. Secondary credentials-If the primary credentials fail, the system determines whether it has secondary credentials for the user. If available, it submits the user's secondary system credentials provided that the conditions described for primary credentials are true.
 4. Last-entered credentials-If the system does not submit secondary credentials or if the credentials fail, it determines whether it has stored credentials for the specified user and a different resource in its cache. If available, the system submits its stored credentials provided the conditions described for primary credentials are true.
 5. User-specified credentials (prompt)-If the system does not submit last-entered credentials or if the credentials fail, it prompts the user to manually enter his credentials in the intermediate sign-in page. If the user selects the "Remember password?" check box, the system caches the user-specified credentials and, if necessary, resubmits them when the user tries to access the same resource again. Note that when the system caches these credentials, it remembers the specific user and resource, even after the user signs out of the device.



The system does not support the multiple credential authentication mechanism described in this section with the SAML SSO mechanisms.

You cannot define a certificate server, or SAML server as a secondary authentication server.

The system supports basic authentication and NTLM challenge/response scheme for HTTP when accessing web applications, but does not support HTTP-based cross-platform authentication via the negotiate protocol.

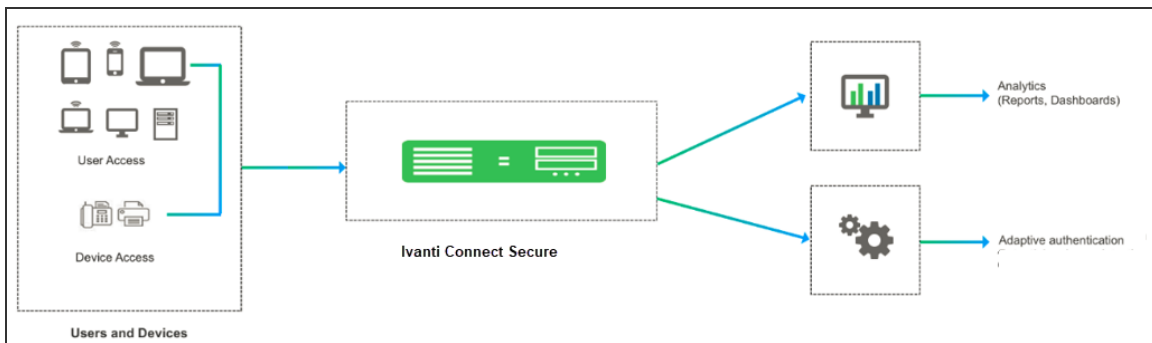
Adaptive Authentication

Overview

Enterprises deal with constant and ever-increasing magnitude of threat vectors, which includes Data Loss Prevention (DLP), Domain Generation Algorithms (DGA) attacks and so on. With changing business requirements and new types of threats, Administrators must understand how users and devices are accessing company's data and services to ensure that the access control policies are up to date. Even after successful authentication, the user's activity should be monitored fully to ensure device compliance.

Behavioral Analytics feature analyzes user's action along with other context data to derive conclusions about any anomalous activities. It provides information/visibility based on real time user or device context thus helping in advanced attack detection and helps in proactive policy-based enforcement.

The following figure depicts Behavioral Analytics:



The Behavioral Analytics feature addresses the following types of anomaly detection:

- User/device is prompted for second level of authentication based on the threat profile determined for the corresponding user/device.

Below are some scenarios where second level of authentication is required:

- User authenticating from new device: This is detected by using the device MAC address.
- User authenticating from new location: Location details are obtained by using the location configurations.

Adaptive Authentication User Flow

1. Users connect to ICS.

2. ICS performs the primary authentication.
3. ICS checks for any anomalies using Behavioral Analytics.
4. ICS prompts for secondary authentication to connect to network to ensure only the valid users accesses the network.
5. User enters the credentials required for secondary authentication.
6. If a user is logging in for the first time or if the user location changes, then ICS performs the secondary authentication and allows/rejects access to the user/device.

Benefits

- ICS monitors the traffic from users and helps in determining the possible anomalous activities such as:
 - If the user is authenticating from a new device / new location.
 - If the device traffic is different from previous instances.
- Data collected as part of Behavior Analytics is stored so that it can be used later for determining the anomalies.

Configurations

- [Summary of Configuration](#)
- [Configuring ICS for Enabling Behavioral Analytics](#)

Summary of Configuration

1. Administrator enables the behavior analytics and configures ICS for Adaptive Authentication
2. Once the anomalies are detected, ICS tags the corresponding user profiles in the data.
3. Administrator configures the role mapping rules to consume these flags and control the access to the corresponding users.
4. Administrator enables the secondary authentication for the users in case they are tagged with anomalies activities to ensure additional level of authentication for security purpose.
5. View the Dashboard and Reports for any detected anomalies.
6. Administrator can also choose to clear the detected anomalies from the Reports page.



Behavior Analytics configuration is synched across the nodes in the cluster (including config-only clusters). However, data collected and analyzed is synched across the nodes but not in case of config-only clusters.

Configuring ICS for Enabling Behavioral Analytics

The Behavioral Analytics package is available by default for detecting the anomalies. If you plan to upgrade to the latest package, it can be downloaded from the [support portal](#).

Configuring ICS for Adaptive Authentication

To enable behavioral analytics:

1. Select **System > Behavioral Analytics > Configuration**.
2. Under Configurations, select **Enable Behavioral Analytics**.
3. For enabling Adaptive Authentication, select **Enable data collection during authentication of devices and users**.

In case you have a Fresh Installation of ICS/IPS, then it will NOT have UEBA package by default with it. Please add the UEBA package before using Adaptive Authentication. In case of Upgrade of ICS/IPS from R7 or earlier to R8 or later, then UEBA package is carried forwarded as is and you can still update it to latest version by uploading new package. You may download latest UEBA package from Support Site (my.pulsesecure.net)

The following figure depicts the Behavioral Analytics Configuration:

Behavioral Analytics > Configuration

Configuration

Configuration

Installed Version: 1.0.3

Last Updated Time: Wed Sep 19 22:31:13 2018

Package: No file chosen

Enable Behavioral Analytics

Enable data collection during authentication of devices and users

4. Navigate to **Administrators > Admin Realms or Users > User Realms**.
5. Under Additional Authentication Server, select **Enable Additional Authentication Server**.
 - Select **Enable adaptive authentication**.
 - Under Authentication #2, select the desired secondary authentication server from the drop-down list.

The following figure depicts the Additional Authentication Server:

Additional Authentication Server

Enable additional authentication server

You can specify an additional authentication server for single sign-on (SSO) purposes. The additional credentials can be specified by the user on the sign-in page (the labels for these inputs are specified by the sign-in page), or they can be pre-defined below, in which case the user will not be prompted for the credential.

Enable adaptive authentication

Note: Adaptive authentication is supported by leveraging the behavioral analytics. Enable behavioral analytics on 'System->Behavioral Analytics->Configuration' for supporting this. Adaptive Authentication is not supported with 'Anonymous' type authentication server selected as authentication server above.

Authentication #2: AA-Local

Username is: specified by user on sign-in page
 predefined as: <USER>

Password is: specified by user on sign-in page
 predefined as: <PASSWORD> Mask static password

End session if authentication against this server fails

Dynamic policy evaluation

Enable dynamic policy evaluation

6. Click **Save Changes**.

Dashboard and Reports

The Behavioral Analytics dashboard provides visibility to many anomalies in the network. It provides visibility of any known, active anomalies, devices with potential malware, IoT devices with anomalous traffic, anomalies location, trend and so on.

To view the Behavioral Analytics dashboard:

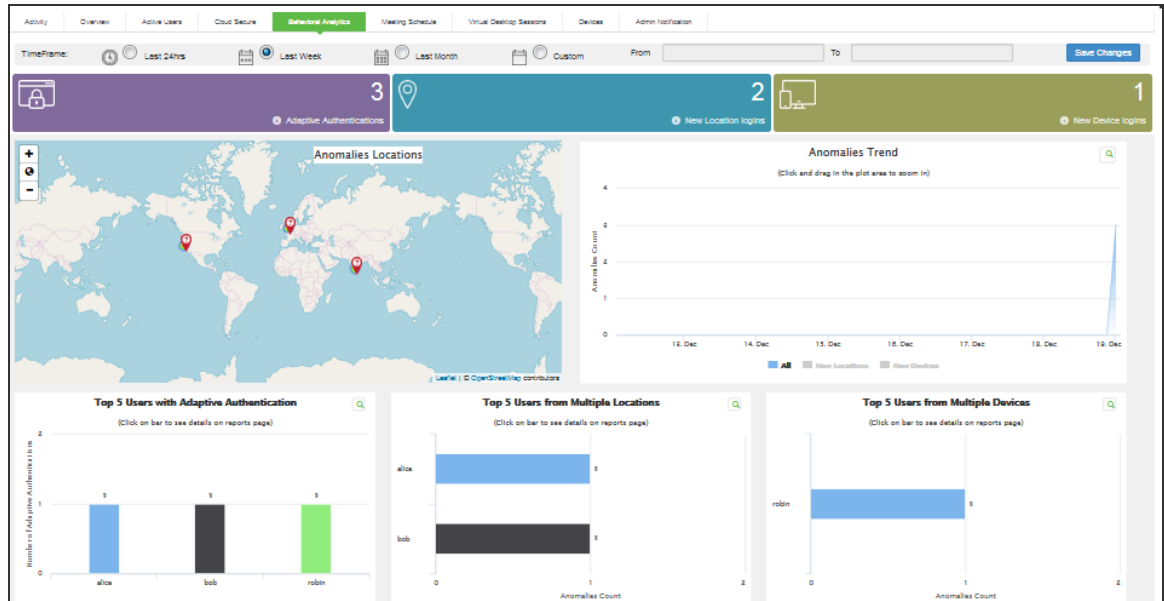
1. Select **System > Status > Behavioral Analytics**.
2. Select the desired timeframe from available options.
3. Click **Save Changes**.

You can also view the drill down reports such as:

- Top 5 Users with Adaptive Authentication


- Top 5 Users from Multiple Locations
- Top 5 Users from Multiple Devices

The following figure depicts the Behavioral Analytics Dashboard Page:



The Reports page is enhanced to view the behavioral analytics related reports. To view the reports, select **System > Reports > Behavioral Analytics**.

The following figure depicts the Behavioral Analytics Reports Page:



Troubleshooting

The event and debug logs can be used for troubleshooting.

The Event logs are generated for the user-related anomalies:

- User authentication from new device/location.

You can use the User Access and Admin Logs in case of any issues. The user access logs are generated whenever there are any user related anomalies such as user logging from new location/new user. The Admin Logs are generated whenever there is a change with Behavioral Analytics options and if there are any changes with respect to application policies.

You can also use **Maintenance > Troubleshooting > Monitoring > Debug Log** for debugging issues.

Synchronizing User Records

About User Record Synchronization

The user record synchronization feature promotes a more consistent user experience by allowing users to retain their bookmarks and individual preferences regardless of which device they log in to.

User record synchronization relies on client-server pairings. The client is the device that users log in to start their remote access. Each client is associated with one primary server and one backup server to store user record data. Clients can be individual appliances or a node within a cluster.

A server in this instance is the device that stores the user data records. Each server can be configured to replicate its user record data to one or more peer servers. Servers are identified by a user-defined logical name. The same logical name can be assigned to more than one authentication server to let you associate authentication servers of different types to the same user. For example, SA1 is an ACE authentication server with user1 who creates a bookmark to www.pulsesecure.net. SA2 is an Active Directory authentication server with the same user1. For the www.pulsesecure.net bookmark to be transferred from SA1/ACE/user1 to SA2/AD/user1 you would assign the logical name "Logical1" to both the ACE server on SA1 and the Active Directory server on SA2.



Cluster VIPs cannot be used as the IP for synchronizing between clients and peers servers.

As long as the logical name is the same, the authentication servers can be different types and different server names and still be associated with a common user. The username must be the same for user record data to be synchronized across the servers. The logical authentication server (LAS) and username combination is what uniquely identifies a user record.

The following user records are synchronized between the client and server:

- Bookmarks
 - Web
 - File
 - Terminal Services
 - JSAM
- Preferences
- Persistent cookies
- Cached passwords

User session data is not synchronized. Persistent cookies, if changed, are synchronized when the user session terminates. All other modifications to the user records are synchronized immediately. User records are stored in cache on the client node prior to being pushed to the servers.

When a user logs in to a client, their data is pulled from the associated server. The pull is performed in the background and does not delay the login process. Users using browsers that do not support JavaScript must manually refresh the index page for updated bookmarks and preferences to appear. For browsers that support JavaScript, users may see a spinning progress indicator and their home page will refresh automatically with updated bookmarks and preferences.

Clients and servers need not be installed with the same Ivanti Connect Secure software version.



User record synchronization uses port 17425. This port number is not configurable. If you are deploying across a firewall, configure your firewall to allow traffic on this port.

To set up user record synchronization, you perform the following tasks:

1. Enable user record synchronization for each participating client and server, identify which ones are the client and which ones are the server and assign a node name to each client and server.
2. Create a shared secret which is used to authenticate the client with the server and the server to its peer servers.
3. On each server, define which clients and peers are allowed to communicate with the server.
4. On each client, define the servers that handle records for each LAS server.
When enabling this feature, you have several options to initialize the user record database. You can:
 - populate the database using user records located in the cache of the client systems.
 - populate the database using user records located in the cache of the server systems.
 - don't pre-populate the database but populate it as users log in and out of the client system.
If you choose the last option, users may not be able to view their saved bookmarks and preferences until the next time they log in, depending on which client they log in to.



User records may not synchronize if the time clocks on the devices are not in sync. We recommend that you use the same NTP server for each node participating in user record synchronization to keep system times accurately adjusted.

The user record synchronization feature will not start automatically after importing a system configuration that has this feature enabled. The workaround is to disable user record synchronization and then enable user record synchronization from the user interface after the configuration import.

Enabling User Record Synchronization

The first step in enabling user record synchronizing is to define the node name and the shared secret used to authenticate between the clients and the servers:

1. Select **System > Configuration > User Record Synchronization > General**.
2. Select the **Enable User Record Synchronization** check box.
3. Enter a unique node name. This name is used when associating a client with a server and is different from the logical name assigned to a server. This node name is also not the same as the cluster node name.
4. Enter the shared secret and confirm it.
The shared secret is the password used to authenticate the client with its servers and the primary server with its peer servers. Use the same shared secret for all clients and servers participating in user record synchronization.
5. Select whether this node is client only or if this node acts as both a client and server.
6. Click **Save Changes**.



If you need to make any changes in this window at a later time, you must deselect the Enable User Record Synchronization check box and click **Save Changes**. Make your edits, select the **Enable User Record Synchronization** check box and save your changes.

Once you enter a name and shared secret, you cannot clear these fields.

Configuring the User Record Synchronization Authentication Server

To set up the authentication server you must define its logical name:

1. Select **Authentication > Auth Servers**.

2. Click the name of the authentication server you want assign a LAS name.
By assigning the authentication server a LAS name, all users that authenticate using the authentication server are associated with this LAS. In this instance, we are referring to the client nodes, not the user record synchronization server nodes.
3. Select the **User Record Synchronization** check box.
4. Enter a logical name to identify this server.
This allows you to share user record data across authentication servers on different devices. By assigning a LAS name to an authentication server, you are implicitly assigning it to all users that authenticate with that auth server. The combination of the user's login name and their LAS name uniquely identifies the user's user record across all user record synchronization servers.
5. Click **Save Changes**.

Configuring the User Record Synchronization Server

To set up the user record synchronization server you must define its peer nodes (optional) and the clients that can access this server.

1. Select **System > Configuration > User Record Synchronization > This Server**.
2. Enter the peer server's node name and IP address, then click Add. To specify more than one peer server, enter each server's node name and IP address individually and click **Add**. There is no limit on the number of peer servers you can add.
Data is replicated from the primary or backup server to its peer servers. If the primary is not available, user data is sent to the backup. User data is then replicated to the peer servers.
3. For each client you want synchronized with this server, enter the client's name and IP address and click **Add**.
Once added, peer servers will have a colored icon next to their name indicating their connection status. Node status is provided to client nodes and LAS mapping servers as well.

Color	Description
Green	Connecting
Yellow	Connecting
Gray	Not connected

Configuring the User Record Synchronization Client

To set up the client, you select the primary and backup server you want this client to synchronize with:

1. Select **System > Configuration > User Record Synchronization > This Client**.
2. Select the **LAS name** you want to synchronize and enter the primary IP of the user record. If you prefer to synchronize with any available server, select **Any LAS**.
3. Enter the primary and optionally a backup server's IP address and then click **Add**. Even if you select Any LAS, you must enter a primary server IP address. Once added, the primary and backup servers have a colored icon next to their name indicating their connection status.

Configuring the User Record Synchronization Database

With the Database tab, you can delete inactive records from the client cache, retrieve statistics about the database, export and import the data and remove user data from the server's database.

To configure the database:

1. Select **System > Configuration > User Record Synchronization > Database**.
2. Select **Auto-delete inactive synchronized user records from the Cache** to remove inactive user records from the cache. This option does not remove user records from the user record database.

When this option is selected, the system performs a check every 15 minutes and deletes user records that meet all of the following criteria:

- There are no active user sessions associated with the user record.
- The user record does not have any custom settings, or the latest version of the user record has been synchronized with the user record database.
- The authentication server associated with the user record database does not have type "local". For example, the "System Local" auth server that is part of the default configuration of the system has a "local" type, so any user records associated with that auth server will not be auto-deleted. However, user records associated with external authentication servers like Radius or LDAP may be deleted, depending on the two prior criteria.

3. Select **Auto-delete user records from the local synchronization database that have been idle for X days** to permanently remove user records from the database located on the server. Enter the number of days user records must be inactive before being deleted. In this instance, "inactive" means that no client has pulled the user record or pushed any modifications to the user record in X days.
4. Click **Retrieve Statistics** to display the number of records in the database. You cannot edit or view records in the database.
5. **Under Export, you export user records to a file. The user records can be exported from the user record database, or from the cache. The exported file can be used to pre-populate the user record database on another node.**
 1. Enter the LAS name of the user records you want to export. If you leave this field blank, all user records are exported. If you enter a LAS name, only user records with the entered LAS name are exported.
 2. To encrypt the exported data, select the **Encrypt the exported data with password** check box and enter the password.
 3. Click **Export** to export the user records from the specified source (cache or database). You will be prompted where to save the file.

6. Under **Import**, you import user records into the synchronization database. The user records can be imported from a file or from the cache. Use the **Import** operation to pre-populate the user record database with user records exported from another node, or with user records from the cache.

1. Click **Browse** to locate the exported file and enter the password if the exported file was encrypted with a password.
 2. Select the **Override Logical Auth Servers in imported user records** with check box to replace the LAS name in each imported user record with the LAS name entered. For example, you change the LAS name, use this option to update the user records with the new name.
 3. Click **Import**.
7. Under **Delete**, specify which user records to permanently remove from the user record database. The options you select apply only to the user record database associated with this server.
- Select **User record with login name and Logical Auth Server** to remove a specific record. The login name and LAS name together uniquely identify a user record. Select this option to remove that record (if it exists).
 - Select **User records with Logical Auth Server** to delete all user records with the specified LAS name.
 - Select **All user records** to permanently remove user records from the database on this node.
8. Click **Delete**.

Scheduling User Record Synchronization Backup

You can configure periodic backups of the user record database. User record synchronization backup can be enabled only on a user record synchronization server.

To back up the user record database:

1. Ensure the system is set up as a user record synchronization server. See **System > Configuration > User Record Synchronization**.

2. Select **Maintenance > Archiving > Archiving Servers**.
3. Select the **Archive User Record Synchronization Database** check box.
4. Specify an archive schedule. Through the options, schedule archives on any combination of weekdays including weekends.



If you schedule an archival operation to occur during the hour that your system switches to Daylight Savings Time (DST) the operation may not occur as scheduled. For example, if your system is set to change to DST at 1:00 a.m. and you have scheduled an archival operation to occur at any time between 1:01 a.m. and 1:59 a.m., the operation is not accomplished, because at 1:00 a.m. the system clock is moved forward to 2:00 a.m. and the system never reaches your archival time for that date.

-
5. Define a specific time when you want the system to archive data or elect to archive data every hour, which produces twenty-four files with unique timestamps.



We recommend you schedule an archival operation during hours when traffic is light in order to minimize its impact to your users. The automatic archiving process compresses files and, if the system is busy, can degrade performance for users. Also, a cluster node may appear unresponsive if the system is busy with traffic and performing archiving simultaneously.

-
6. Provide a password if you want to encrypt user record synchronization database archives with a password (optional).
 7. Click **Save Changes**.

Host Checker



From 9.1R15 onwards, support for Solaris, and features Cache Cleaner and Statement of Health are deprecated. Ensure you remove all related configurations before upgrading to 9.1R15.

Upgrade may fail if all configurations are not removed. For more information refer [KB45044](#).

Host checker is a client-side agent that performs endpoint health and security checks for hosts that attempt to connect to a Ivanti Connect Secure device. It supports two types of rules within a policy; predefined and custom. The pre-defined inspection capabilities consist of health and security checks including antivirus versions, antispysware, OS versions, hard disk encryption status and patch checks. The pre-defined rules are provided by OPSWAT and it uses the ESAP plug-in for pre-defined checks.

Custom rules allow admin to define checks to collect system health using Integrity message collector (IMC) and evaluate using Integrity message verifier (IMV) of TNC framework. The custom rules are created by the admin to include inspection checks such as absence or presence of specific file, certificate checks, TCP ports, processes, registry key settings, NetBIOS name, MAC addresses or certificate of the client machine and third-party inspection methods (custom DLLs).

You can invoke Host Checker at the role level, or the realm level to specify access requirements for endpoints attempting to authenticate.

For details about the configuration, refer to *Host Checker Configuration Guide* available on the <https://www.ivanti.com/support/product-documentation> site.

Sign-In Policies

About Sign-In Policies

Sign-in policies define the URLs that users and administrators use to access the device and the sign-in pages that they see. The system has two types of sign-in policies—one for users and one for administrators. When configuring sign-in policies, you associate realms, sign-in pages, and URLs.

For example, in order to allow all users to sign in to the device, you must add all user authentication realms to the user sign-in policy. You may also choose to modify the standard URL that the end-users use to access the system and the sign-in page that they see. Or, if you have the proper license, you can create multiple user sign-in policies, enabling different users to sign into different URLs and pages.

You can create multiple sign-in policies, associating different sign-in pages with different URLs. When configuring a sign-in policy, you must associate it with a realm or realms. Then, only members of the specified authentication realm(s) may sign in using the URL defined in the policy. Within the sign-in policy, you may also define different sign-in pages to associate with different URLs.

For example, you can create sign-in policies that specify:

- Members of the "Partners" realm can sign in to the device using the URLs: partner1.yourcompany.com and partner2.yourcompany.com. Users who sign into the first URL see the "partners1" sign-in page; users who sign into the second URL see the "partners2" sign-in page.
- Members of the "Local" and "Remote" realms can sign into the device using the URL: employees.yourcompany.com. When they do, they see the "Employees" sign-in page.
- Members of the "Admin Users" realm can sign into the device using the URL: access.yourcompany.com/super. When they do, they see the "Administrators" sign-in page.

When defining sign-in policies, you may use different hostnames (such as partners.yourcompany.com and employees.yourcompany.com) or different paths (such as yourcompany.com/partners and yourcompany.com/employees) to differentiate between URLs.



- If a user attempts to sign in while there is another active user session with the same sign-in credentials, the system displays a warning page showing the IP address of the existing session and two buttons: Continue and Cancel. By clicking the Cancel button, the user terminates the current sign-in process and redirects the user back to the Sign-in page. By clicking the Continue button, the system creates the new user session and terminates the existing session.
- When enabling multiple sign-in URLs, note that in some cases the system must use cookies on the user's machine to determine which sign-in URL and corresponding sign-in page to display to the user. The system creates these cookies when the user signs into the device. (When a user signs into the device, the system responds with a cookie that includes the sign-in domain of the URL. The system then attaches this cookie to every system request the user makes.) Generally, these cookies ensure that the system displays the correct sign-in URL and page to the user. For example, if a user signs into the device using the URL <http://yourcompany.net/employees> and then her session times out, the system uses the cookie to determine that it must display the <http://yourcompany.net/employees> sign-in URL and corresponding page to the user when she requests another system resource.

However, in isolated cases, the cookie on the user's machine may not match the resource she is trying to access. The user may sign into one URL and then try to access a resource that is protected by a different URL. In this case, the system displays the sign-in URL and corresponding sign-in page that the user signed into most recently. For example, a user may sign into the device using the sign-in URL <http://yourcompany.net/employees>. Then she may try to access the system resource using a link on an external server, such as <https://yourcompany.net/partners/dana/term/winlaunchterm.cgi?host=<termsrvIP>>. Or, she may try to open a bookmark that she created during a different session, such as <https://yourcompany.net/partners/,DanaInfo=.awxyBmszGr3xt1r5O3v.,SSO=U+>. In these cases, the system would display the <http://yourcompany.net/employees> sign-in URL and page to the user, rather than the sign-in URL or page that is associated with the external link or saved bookmark that she is trying to access.

Sign-in policies and pages are an integral part of the access management framework, and therefore are available in all Ivanti Connect Secure products.

Task Summary: Configuring Sign In Pages

To configure sign-in policies, you must:

1. Create an authentication realm through the **Administrators > Admin Realms or the Users > User Realms** page of the admin console.
2. (Optional) Modify an existing sign-in page or create a new one using options in the **Authentication > Signing In > Sign-in Pages** page of the admin console.
3. (Optional) Modify an existing sign-in page or create a new one using options in the **Authentication > Signing In > Sign-in Pages** page of the admin console.
4. Specify a sign-in policy that associates a realm, sign-in URL, and sign-in page using settings in the **Authentication > Signing In > Sign-in Policies** page of the admin console.
5. If you differentiate between URLs using hostnames, you must associate each hostname with its own certificate or upload a wildcard certificate into the system using options in the **System > Configuration > Certificates > Device Certificates** page.

About Configuring Sign In Policies

User sign-in policies also determine the realm(s) that users and administrators can access.

Depending on whether a sign-in policy is for endpoints (users) or administrators, the configuration options are different. For users, different authentication protocol sets can be configured, and realm selection is based on the authentication method that is associated with the realm.

Configuring User Sign In Policies

To create or configure user sign-in policies:

1. In the admin console, select **Authentication > Signing In > Sign-in Policies**.
2. To create a new sign-in policy, click **New URL**. Or, to edit an existing policy, click a URL in the Administrator URLs or User URLs column.
3. Select **Users** or **Administrators** to specify which type of user can sign in using the access policy.
4. In the Sign-in URL field, enter the URL that you want to associate with the policy. Use the format `<host>/<path>` where `<host>` is the hostname of the device, and `<path>` is any string you want users to enter. For example: `partner1.yourcompany.com/outside`. To specify multiple hosts, use the `*` wildcard character.

To specify that all administrator URLs should use the sign-in page, enter `*/admin`.

- You may only use wildcard characters (*) in the beginning of the hostname portion of the URL. The system does not recognize wildcards in the URL path.
 - SAML authentication does not support sign-in URLs that contain multiple realms. Instead, map each sign-in URL to a single realm.
5. (optional) Enter a Description for the policy.
 6. **From the Sign-in Page list**, select the sign-in page that you want to associate with the policy. You may select the default page that comes with the system, a variation of the standard sign-in page, or a custom page that you create using the customizable UI feature.
 7. Under Authentication realm, specify which realm(s) map to the policy, and how users and administrators should pick from amongst realms. If you select:
 - **User types the realm name**-The system maps the sign-in policy to all authentication realms, but does not provide a list of realms from which the user or administrator can choose. Instead, the user or administrator must manually enter his realm name into the sign-in page.
 - **User picks from a list of authentication realms**-The system only maps the sign-in policy to the authentication realms that you choose. The system presents this list of realms to the user or administrator when he signs-in to a device and allows him to choose a realm from the list. (Note that the system does not display a drop-down list of authentication realms if the URL is only mapped to one realm. Instead, it automatically uses the realm you specify.)
 8. Click **Save Changes**.

Enabling and Disabling Sign-In Policies

To enable and disable sign-in policies:

1. In the admin console, choose **Authentication > Signing In > Sign-in Policies**.
2. To enable or disable:
 - **An individual policy**-Select the check box next to the policy that you want to change, and then click Enable or Disable.

If you select this option, all user sessions are immediately terminated. If this device is part of a cluster, all user sessions across all nodes in the cluster are immediately terminated.

3. Click **Save Changes**.

Specifying the Order in Which Sign-In Policies are Evaluated

The system evaluates sign-in policies in the same order that you list them on the Sign-in Policies page. When it finds a URL that matches exactly, it stops evaluating and presents the appropriate sign-in page to the administrator or user. For example, you may define two administrator sign-in policies with two different URLs:

1. The first policy uses the URL `*/admin` and maps to the default administrator sign-in page.
2. The second policy uses the URL `yourcompany.com/admin` and maps to a custom administrator sign-in page.

If you list the policies in this order on the Sign-in Policies page, the system never evaluates or uses the second policy because the first URL encompasses the second. Even if an administrator signs in using the `yourcompany.com/admin` URL, the system displays the default administrator sign-in page. If you list the policies in the opposite order, however, the system displays the custom administrator sign-in page to those administrators who access the system using the `yourcompany.com/admin` URL.

Note that the system only accepts wildcard characters in the hostname section of the URL and matches URLs based on the exact path. For example, you may define two administrator sign-in policies with two different URL paths:

1. The first policy uses the URL `*/marketing` and maps to a custom sign-in page for the entire Marketing Department.
2. The second policy uses the URL `*/marketing/joe` and maps to a custom sign-in page designed exclusively for Joe in the Marketing Department.

If you list the policies in this order on the Sign-in Policies page, the system displays Joe's custom sign-in page to him when he uses the `yourcompany.com/marketing/joe` URL to access the device. He does not see the Marketing sign-in page, even though it is listed and evaluated first, because the path portion of his URL does not exactly match the URL defined in the first policy.

To change the order in which administrator sign-in policies are evaluated:

1. In the admin console, choose **Authentication > Signing In > Sign-in Policies**.
2. Select a sign-in policy in the Administrator URLs, User URLs list.
3. Click the up and down arrows to change the selected policy's placement in the list.
4. Click **Save Changes**.

Configuring Fallback Authentication Server

In case the remote authentication server is not reachable, a local authentication server acts as a fallback.

This feature is currently available in the AD server and LDAP servers.

The administrator can use a randomly generated URL to sign in. The local authentication server supports the randomly generated URL to sign in within 10 minutes of generation. The URL is randomly generated and provided on the Admin login screen. Disabling the URL is optional.

To set a fallback URL:

1. In the Admin UI choose **Authentication > Signing In > Sign-in Policies**.
2. Create a new Administrator URL in **Authentication > Signing In > Sign-in Policies > Admin URL configuration** with localauth and disable the same.
3. Under the **Configure Fallback URL** select the option for fallback.
4. Select the **Backup URL** from the dropdown created in step 3 above and **Save changes**.

About Sign-In Notifications

With sign-in notifications, you can create and configure detailed notification messages that appear for Pulse clients and for agentless access endpoints when the user attempts to sign in. For example, you could configure a notification message that explains terms of use, company-specific policies, a welcome page, an end user license agreement (EULA), or a message of the day (MOTD).

For a browser-based (agentless) login, the notification message appears in a separate page either before (pre-auth) or after (post-auth) user authentication during the sign-in process. For a Pulse client login, the notification messages appear in a Ivanti message box. The user is expected to read the content of the sign-in notification message and acknowledge by clicking a Proceed button. The user may indicate disagreement by clicking a Decline button, which ends the login attempt.

You can configure a sign-in policy to use a sign-in notification either as pre-auth or post-auth (or both). In the case of post-auth configuration, you can either use a common message for all roles or use separate messages for each role.

You can create a multi-language sign-in notification package that relies on the language setting of the endpoint. You can customize the sign-in notification page appearance for browser-based logins by modifying the related fields in a sign-in page in the Admin UI or by using a custom sign-in page.

- Sign-in notifications are supported on Windows, Mac, and for browser-based access on mobile devices. However, sign-in notifications might not work well with all mobile devices due to device limitations.
- Sign-in notifications (including uploaded packages) are included in XML exports.
- If a Ivanti session is resumed or extended, the pre-auth notification message is not shown again. However, if the user switches roles when resuming a session, and that role change results in a new notification, Ivanti displays the message. You can configure the post-auth message to be skipped if it has already been seen. If the post-auth message is not marked to be skipped, then it always appears.

Configuring and Implementing Sign-in Notifications

Sign-in notifications appear for Pulse client and for browser-based logins when the user attempts to sign in.

To configure and implement sign-in notifications:

1. In the admin console, select **Authentication > Signing In > Sign-in Notifications**.
2. Click **New Notification**.
3. Specify a Name for the notification. This name appears in the sign-in policies page, and in the UI Options page for a selected role.
4. Select **Text or Package** in the Type box.
 - If you select Text, type the desired sign-in notification message, or copy and paste the relevant text into the Text field.
 - If you select Package, click the **Browse** button and navigate to a previously prepared .zip file. A package is typically used to provide different language versions of the notification message.
 - The zip file should include a default.txt file and one or more <language>.txt files (Example: en.txt).
 - Language-abbreviations should be strings that can appear in Accept-Language header of an HTTP request. For example:

1. Upload a zip file containing files with name format:
<language-abbreviation>.txt (Example: en.txt).
 2. Include 'default.txt' and one file for each language you want to support.
 3. Language-abbreviations should be strings that can appear in Accept-Language header of an HTTP request.
- The character encoding supported is UTF-8.



When you create a zip file, do not add the folder containing the files, but add the files directly.

5. Click **Save Changes**.

To enable sign-in notifications:

1. In the admin console, click **Authentication > Signing In > Sign-in Policies**.
2. Select an existing URL or create a new URL.
3. **Under Configure Sign-in Notifications**, select the check box for Pre-Auth Sign-in Notification, Post-Auth Sign-in Notification, or both.
 - **After Pre-Auth Sign-in Notification**, select a previously configured sign-in notification from the drop-down menu.
 - **After Post-Auth Sign-in Notification**, select the option for Use a common Sign-in Notification for all roles or Use the Sign-in Notification associated to the assigned role.
 - If you select Use a common Sign-in Notification for all roles, select a previously configured sign-in notification from the drop-down menu.
 - If you select Use the Sign-in Notification associated to the assigned role, the sign-in notification configured for the assigned role will be used.
 - Prevent the Post-Auth sign-in notification from being displayed to users who have seen it before, by selecting the Skip if already shown check box. (This is only a hint to the system and might not be honored in all environments.)
4. Click **Save Changes**.

5. You can customize the appearance of the sign-in notification message by selecting **Authentication > Signing In > Sign-in Pages** and creating a sign-in page or using an existing page.
6. Under Sign-in Notification appearance, customize UI options for Pre-Auth Notifications and Post-Auth Notifications by changing the following items:
 - For Notification Title enter the text that appears at the top of the sign-in notification page.
 - In the Proceed Button box, enter the text for the button that the user clicks to proceed with the sign-in.
 - This text applies to browser-based logins only. A Pulse client login always displays Proceed.
 - Optionally, clear the check box for Display "Decline" Button. If this box is not checked, the user does not have the option to decline.
 - In the Decline Button box, enter the text for the button that the user clicks to decline.
 - This text applies to browser-based logins only. A Pulse client login always displays Decline.
 - In the Message on Decline box, enter the text that you would like to appear when a user clicks the Decline button.
7. Click **Save Changes**.



- When Console Protection is enabled for the ICS console, the Sign-In Notification configured for /admin Sign-In URL is displayed on the ICS Console. However, if the Sign-In Notification is loaded from a package, a default banner message is displayed on the console.
- If you enabled Use the Sign-in Notification associated to the assigned role you must complete the implementation by selecting the sign-in notification on the Users > User Roles > Role Name > General > UI Options page or Administrators > Admin Roles > Role Name > General > UI Options page, as applicable.

If more than one role is available to a user, the sign-in notification associated with the first role assigned is displayed.

8. Add the sign-in page in which you have customized the sign-in notification appearance to the sign-in policy.

Defining Authorization-Only Access Policies

Authorization-only access is similar to a reverse proxy. Typically, a reverse proxy is a proxy server that is installed in front of web servers. All connections coming from the Internet addressed to one of the web servers are routed through the proxy server, which may either deal with the request itself or pass the request wholly or partially to the main web server.

With an authorization-only access, you select a user role. The system then acts as reverse proxy server and performs authorization against the server for each request.

For example, the authorization-only access feature satisfies the following business needs:

- If you have a third-party AAA policy management server, the system acts as an authorization-only agent.
- If your user sessions are managed by a third-part session management system, there is no need to duplicate the user session management in the system.

With authorization-only access, there is no SSO from the system. SSO is controlled by your third-party AAA infrastructure.

To create or configure authorization-only access policies:

1. In the admin console, choose **Authentication > Signing In > Sign-in Policies**.
2. To create a new authorization only access policy, click **New URL** and select authorization only access. Or, to edit an existing policy, click a URL in the Virtual Hostname column.
3. In the Virtual Hostname field, enter the name that maps to the system's IP address. The name must be unique among all virtual hostnames used in pass-through proxy's hostname mode. The hostname is used to access backend application entered in the Backend URL field. Do not include the protocol (for example, http:) in this field.

For example, if the virtual hostname is myapp.ivehostname.com, and the backend URL is http://www.xyz.com:8080/, a request to https://myapp.ivehostname.com/test1 via the system is converted to a request to http://www.xyz.com:8080/test1. The response of the converted request is sent to the original requesting web browser.

4. In the Backend URL field, enter the URL for the remote server. You must specify the protocol, hostname and port of the server. For example, http://www.mydomain.com:8080/*.

When requests match the hostname in the Virtual Hostname field, the request is transformed to the URL specified in the Backend URL field. The client is directed to the backend URL unaware of the redirect.

5. (optional) Enter a Description for this policy.
6. Select the server name or No Authorization from the Authorization Server drop-down menu. If you select a server, ensure that the front-end server provides the SMSESSION cookie otherwise you will receive an error.
7. Select a user role from the Role Option drop-down menu.

Only the following user role options are applicable for authorization-only access.

- Allow browsing un-trusted SSL web sites (Users > User Roles > RoleName > Web > Options > View advanced options)
- HTTP Connection Timeout (Users > User Roles > RoleName > Web > Options > View advanced options)
- Source IP restrictions (Users > User Roles > RoleName > General > Restrictions)
- Browser restrictions (Users > User Roles > RoleName > General > Restrictions)

Ensure the user role you select has an associated Web Access policy.

8. Select the Allow ActiveSync Traffic only option to perform a basic of validation of the HTTP header to ensure the request is consistent with ActiveSync protocol. If you select this option only ActiveSync protocol requests can be processed. If validation fails, a message is created in the user's event log. If you do not select this option, both ActiveSync and non-ActiveSync requests are processed.
9. Select the Kerberos Constrained Delegation Label option to configure a KCD policy for Active Sync. This would list the existing configured Constrained Delegation labels. Selecting any one of the valid Constrained Delegation labels would force to use KCD for the Exchange Active Sync traffic. Also, this option is applicable only for Active Sync traffic.

This option also has the following dependencies:

- Enforce client certificate requirement on virtual ports which are used for Active Sync.
- Appropriate CA certificate should be imported under Trusted Client CAs.
- The role configured to use for Active Sync should be configured to have Certificate Restrictions to Only allow users with a client-side certificate signed by Certification Authority to sign in.
- Appropriate Constrained Delegation policy should be configured. Please refer to the section "Constrained Delegation" under configuring SSO policies



External configurations should be appropriately configured to support Constrained Delegation SSO; Exchange server should be configured to allow Kerberos authentication, i.e., **Windows Authentication**.

10. If **Kerberos Constrained Delegation Label** policy is chosen, enter the appropriate Username Template from certificate attributes.
11. Click **Save Changes** to save your edits.

The System Status Overview page displays the number of current active concurrent connections and a histogram of the active concurrent connections (Authorization Only Access Active Connections plot in the Concurrent SSL Connections graph).

Configuring Sign-In Pages

A sign-in page defines the customized properties in the end-user's welcome page such as the welcome text, help text, logo, header, and footer. The system allows you to create two types of sign-in pages to present to users and administrators:

- Standard sign-in pages-Standard sign-in pages are produced by Ivanti and are included with all versions of the Ivanti Connect Secure software. You can modify standard sign-in pages through the **Authentication > Signing In > Sign-in Pages** tab of the admin console.
- Customized sign-in pages-Customized sign-in pages are THTML pages that you produce using the Template Toolkit and upload to the system in the form of an archived ZIP file. The customized sign-in pages feature enables you to use your own pages rather than having to modify the sign-in page included with the system.

Configuring Standard Sign-In Pages

Standard sign-in pages that come with the system include:

- Default Sign-In Page-the system displays this page to users when they sign into the device.

You can modify the default sign-in page that the system displays to users when they sign into the device. You can also create new standard sign-in pages that contain custom text, logo, colors, and error message text using settings in the Authentication > Signing In > Sign-in Pages tab of the admin console.

To create or modify a standard sign-in page:

1. In the admin console, select **Authentication > Signing In > Sign-in Pages**.

2. If you are:
 - **Creating a new page**-Click **New Page**.
 - **Modifying an existing page**-Select the link corresponding to the page you want to modify.
3. (New pages only) Under Page Type, specify whether this is an administrator/user access page.
4. Enter a name to identify the page.
5. In the Custom text section, revise the default text used for the various screen labels as desired. When adding text to the Instructions field, note that you may format text and add links using the following HTML tags: `<i>`, ``, `
`, ``, and `<a href>`. However, the system does not rewrite links on the sign-in page (since the user has not yet authenticated), so you should only point to external sites. Links to sites behind a firewall will fail.

If you use unsupported HTML tags in your custom message, the system may display the end-user's home page incorrectly.

6. In the Header appearance section, specify a custom logo image file for the header and a different header color.
7. In the Custom error messages section, revise the default text that is displayed to users if they encounter certificate errors.

You can include `<<host>>`, `<<port>>`, `<<protocol>>`, and `<<request>>` variables and user attribute variables, such as `<<userAttr.cn>>` in the custom error messages. Note that these variables must follow the format `<variable>` to distinguish them from HTML tags which have the format `<tag>`.

8. To provide custom help or additional instructions for your users, select Show Help button, enter a label to display on the button, and specify an HTML file to upload to the system. Note that the system does not display images and other content referenced in this HTML page.
9. Click Save Changes. The changes take effect immediately, but users with active sessions might need to refresh their Web browsers.
10. Click Restore Factory Defaults to reset the sign-in page, the user home page, and admin console appearance.

Configuring Custom Sign-In Pages

To upload custom sign-in Pages into ICS:

1. Download new "Sample Custom Page" from new Admin UI after login as Admin.
(Authentication -> Sign-In Pages -> Upload Custom Pages -> Click on "Sample" It will download the Sample Folder as ZIP & save it on Local disk)
2. Copy the following files after unzip the folder (locally saved in previous step):
 - Logout.thtml
 - PleaseWait.thtml
3. Open pre-downloaded Sample Custom Sign-in folder as unzipped and replace all those files here.
4. Now Select all the files and create *.ZIP file to uploading custom sign-in page on latest build.
5. Log into ICS as admin which is running on latest build and follow the steps to upload new Custom Sign-In Page- In new Admin UI (Authentication -> Sign-In Pages -> Upload Custom Pages -> Put the name of Custom Sign-In Page -> Click on Browse "Button" and select previously saved *.ZIP file from local storage in step-4 -> Now click on "Upload Custom Pages" After successfully Upload finally click on "Save Changes")
6. Once all the above steps are successful, we can see a New Sign-In Pages has been added under Authentication -> Sign-In Pages.

Preventing Sign-In URL Tampering

This feature ensures that the hostname of the current URL matches the one that is associated with the internal id embedded in URL. This feature is not enabled by default and has to be enabled by using XML import.

To enable this feature, use the following HTML tags:

```
<system>
<configuration>
<security>
<signin-url-check>mitigate-url-tamper<signin-url-check>
<security>
<configuration>
<system>
```

Resource Profiles

Resource Profiles



From 9.1R15 onwards, support for Telnet/SSH, Citrix web interface/JICA, Citrix Storefront(below 3.1 version), Microsoft OWA 2000, 2003,2007, Lotus inotes 5, 6, 6.5 are deprecated. Ensure you remove all related configurations before upgrading to 9.1R15. **Upgrade may fail if all configurations are not removed.** For more information refer [KB45044](#).

A resource profile contains all of the resource policies, role assignments, and end-user bookmarks required to provide access to an individual resource. Resource profiles simplify resource configuration by consolidating the relevant settings for an individual resource into a single page within the admin console.

The system comes with two types of resource profiles:

- Standard resource profiles enable you to configure settings for a variety of resource types, such as web sites, client/server applications, directory servers, and terminal servers. When you use this method, you choose a profile type that corresponds to your individual resource and then provide details about the resource.
- Resource profile templates enable you to configure settings for specific applications. When you use this method, you choose a specific application (such as the Citrix NFuse version 4.0). Then, the system pre-populates a variety of values for you based on your chosen application and prompts you to configure additional settings as necessary.

Resource profiles are an integral part of the access management framework, and therefore are available on all Ivanti Connect Secure products. However, you can only access resource profile types that correspond to your licensed features.

To create resource profiles, you:

- Create user roles through the **Users > User Roles** page of the admin console.
- Create resource profiles through the **Users > Resource Profiles** page of the admin console. When creating the resource profile, specify the resource, create autopolicies, associate the profile with user roles, and create bookmarks as necessary.

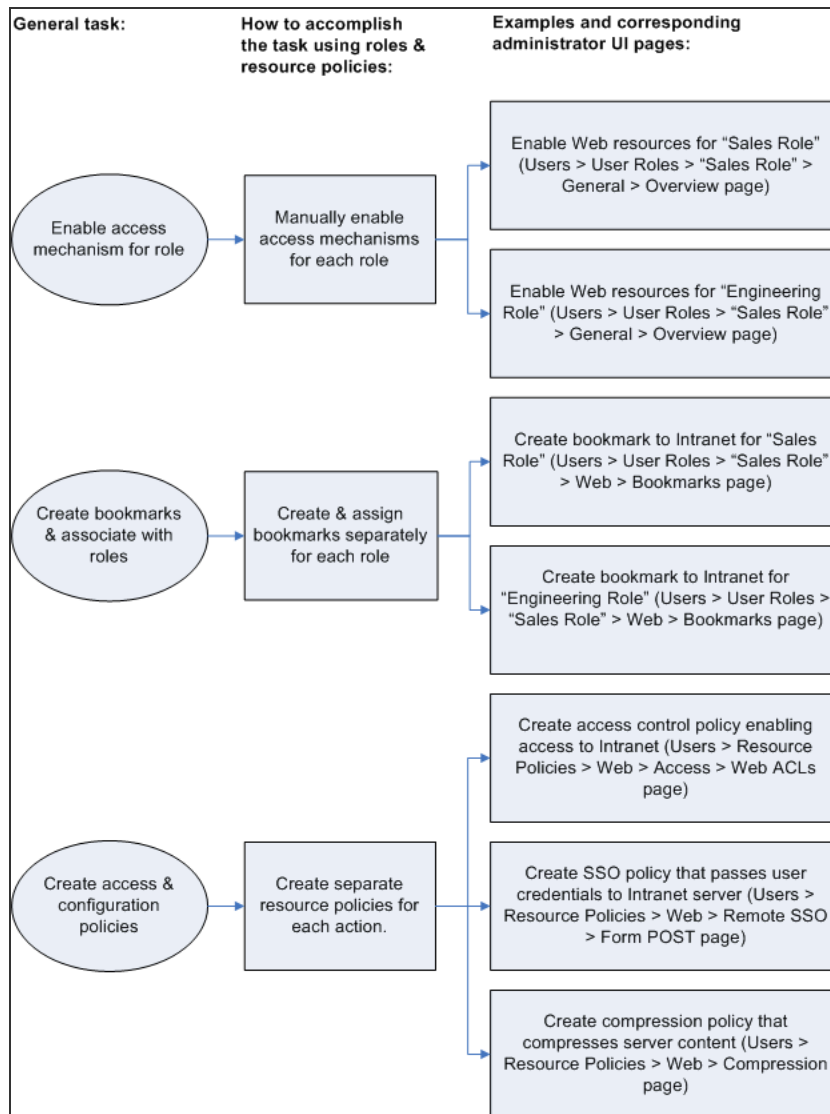
Resource Profile Components

Resource profiles contain the following components:

- **Resources** - When you are defining a resource profile, you must specify the individual resource that you want to configure (such as your company Intranet site or a Lotus Notes application). All other major settings within the profile branch from this resource. You can configure a variety of resource types, including web sites, client/server applications, directory servers, and terminal servers.
- **Autopolicies** - When you are defining a resource profile, you generally create autopolicies that establish the access requirements and other settings for the specified resource. The most common type of autopolicy enables access to the primary resource defined in the profile. Other policy types (such as compression and caching autopolicies) "fine-tune" how the system handles the data that it passes to and from the specified resource.
- **Roles** - When you are defining a resource profile, you generally associate the profile with user roles. The specified roles then inherit the autopolicies and (optionally) the bookmarks defined in the resource profile.
- **Bookmarks** - When you are defining a resource profile, you may optionally create a bookmark that links to the profile's primary resource (such as your company intranet's main page). You can also create additional bookmarks that link to various sites within the resource's domain (such as the Sales and Marketing intranet pages). The system displays these bookmarks to users who are assigned to the user roles that you specify.

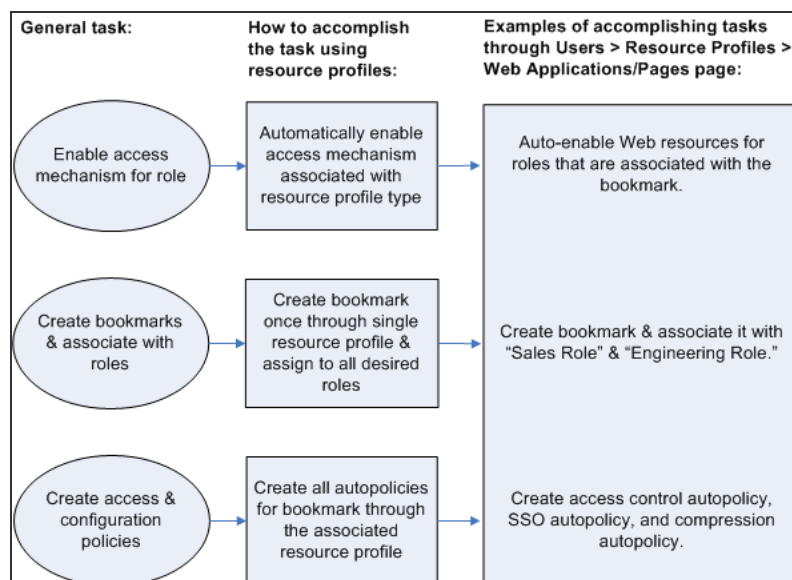
The following figure shows how to configure resources using roles and resource policies. Note that to enable a bookmark for multiple user roles, you must manually re-create the bookmark and enable the appropriate access mechanism for each role. You must also use a variety of pages in the administrator console to create associated resource policies enabling access to the resource and other configuration options.

The following figure depicts Using Roles and Resource Policies to Configure Resources:



The following figure shows how to configure resources using resource profiles. Note that you can create a bookmark, associate it with multiple user roles, and create the associated autopolicies enabling access to the resource and other configuration options through a single section in the administrator console. Also note that the system automatically enables the appropriate access mechanism to the roles to which you assign the bookmark.

The following figure depicts Using Resource Profiles to Configure Resources:



Defining Resource Profile Resources

When you are defining a resource profile, you must specify the individual resource that you want to configure. [Table](#) shows the dependency between the type of profile you choose and the resource you want to configure.

The following table lists the Resource Profile Types and Configuration Information

Use this type of resource profile	To configure this type of resource
Web application/pages	URLs to Web applications, Web servers, and Web pages; Java applets that are stored on third party servers.
Host Java applet	Java applets that you upload directly to the device.
SAM client application	Client/server applications
PSAM destination	Destination networks or servers
Terminal Services	Windows and Citrix terminal servers



You cannot configure applications through VPN Tunneling using resource profiles. Instead, you must use roles and resource policies.

When defining resources, you can use Connect Secure variables, such as <user> to dynamically link users to the correct resources. For instance, you can specify the following Web resource in order to direct users to their own individual intranet pages:

http://yourcompany.intranet/<user>

If the resource field of two different resource profiles are identical and both resource profiles are mapped to the same role, a user might view a resource policy from one profile and a resource policy from the other resource profile. For example, consider the following:

- **Resource Profile #1:Resource Profile Name: Intranet**
- **Resource Profile resource: http://intranet.company.com**
- **Resource Profile Web ACL: http://intranet.company.com/sales/***
- **Mapped to Role: Sales**
- **Resource Profile #2:**
- **Resource Profile Name: Intranet for Sales**
- **Resource Profile resource: http://intranet.company.com**
- **Resource Profile Web ACL: http://intranet.company.com/sales/docs/***

The end user that maps into the Sales role might see a bookmark name Intranet for Sales, but the Web ACL enforcement will be *http://intranet.company.com/sales/**.

This type of configuration is not supported.

Defining Resource Profile Autopolicies

When you are defining a resource profile, you generally create autopolicies that establish the access requirements and other settings for the specified resource. The most common type of autopolicy enables access to the primary resource defined in the profile. Other policy types (such as compression and caching autopolicies) "fine-tune" how the system handles the data that it passes to and from the specified resource.

When creating resource profiles, the system only displays those autopolicies that are relevant to the resource profile type. For instance, you may choose to enable access to a client/server application through a PSAM resource profile. When you do, the system displays autopolicies that you can use to enable access to the specified application's server. On the other hand, the system does not display Java access control autopolicies, since Java settings do not apply to PSAM.



When defining access policies, you must explicitly list each hostname address. The policy checking system does not append or use the default domain or search domains in the system network settings.

Additionally, the system consolidates all of the relevant autopolicy options in a single page of the user interface, enabling you to understand all of the configuration possibilities and requirements for any given resource type.



Access control autopolicies are generally based on the primary resource that you define in the resource profile. If you change the profile's primary resource, however, the system does not necessarily update the corresponding autopolicies. You should re-evaluate your autopolicies after changing the profile's primary resource.

For administrators who are accustomed to using a pre-5.3 version of the Ivanti Connect Secure product, note that autopolicies are resource policies. The system allows you to sort and order autopolicies along with standard resource policies in the Users > Resource Policies pages of the admin console. However, the system does not allow you to access more detailed configuration options for autopolicies through this section of the admin console. Instead, if you want to change the configuration of an autopolicy, you must access it through the appropriate resource profile.

For administrators who are accustomed to using a pre-5.3 version of the Ivanti Connect Secure product, note that you can also automatically create resource policies by enabling the Auto-allow option at the role level. However, note that we recommend that you use autopolicies instead, since they directly correspond to the resource you are configuring rather than all resources of a particular type. (You may also choose to enable the Auto-allow option for a role-level feature and create autopolicies for resources of the same type. When you do, the system creates policies for both and displays them in the appropriate resource policies page of the admin console.)

Defining Resource Profile Roles

Within a resource profile, you can assign user roles to the profile. For instance, you might create a resource profile specifying that members of the "Customers" role can access your company's Support Center, while members of the "Evaluators" role cannot. When you assign user roles to a resource profile, the roles inherit all of the autopolicies and bookmarks defined in the resource profile.

Since the resource profile framework does not include options for creating roles, you must create user roles before you can assign them to resource profiles. However, the resource profile framework does include some user role configuration options. For instance, if you assign a user role to a Web resource profile, but you have not enabled Web rewriting for the role, the system automatically enables it for you.



Note that you can assign roles to a resource profile through the role framework as well as the resource profile framework.

Defining Resource Profile Bookmarks

When you create a resource profile, the system generally creates a bookmark that links to the profile's primary resource (such as your company intranet's main page). Optionally, you may also create additional bookmarks that link to various sites within the primary resource's domain (such as the Sales and Marketing intranet pages). When you create these bookmarks, you can assign them to user roles, thereby controlling which bookmarks users see when they sign into the end-user console.



PSAM and JSAM resource profiles do not include bookmarks, since the system cannot launch the applications specified in the resource profiles.

For example, you may create a resource profile that controls access to your company intranet. Within the profile, you may specify:

- Resource profile name: Your Intranet
- Primary resource: `http://intranet.com`
- Web access control autopolicy: Allow access to `http://intranet.com:80/*`
- Roles: Sales, Engineering

When you create this policy, the system automatically creates a bookmark called "Your Intranet" enabling access to `http://intranet.com` and displays the bookmark to members of the Sales and Engineering roles.

You may then choose to create the following additional bookmarks to associate with the resource profile:

- Sales Intranet" bookmark: Creates a link to the `http://intranet.com/sales` page and displays the link to members of the Sales role.
- Engineering Intranet" bookmark: Creates a link to the `http://intranet.com/engineering` page and displays the link to members of the Engineering role.

When configuring bookmarks, note that:



- You can only assign bookmarks to roles that you have already associated with the resource profile-not all of the roles defined on the system. To change the list of roles associated with the resource profile, use settings in its Roles tab.
 - Bookmarks simply control which links the system displays to users-not which resources the users can access. For instance, in the example used above, a member of the Sales role would not see a link to the Engineering Intranet page, but he could access it by entering <http://intranet.com/engineering> his Web browser's address bar. Similarly, if you delete a bookmark, users can still access the resource defined in the profile.
 - The system allows you to create multiple bookmarks to the same resource. If you assign duplicate bookmarks to the same user role, however, the system Service only displays one of them to the users.
 - Bookmarks link to the primary resource that you define in the resource profile (or a sub-directory of the primary resource). If you change the profile's primary resource, the system updates the corresponding bookmarks accordingly.
-

Resource Profile Templates

Resource profile templates enable you to configure settings for specific applications. When you use this method, you choose a specific application. Then, the system pre-populates a variety of values for you based on your chosen application and prompts you to configure additional settings as necessary.

Currently, the system includes templates for the following third-party applications:

- **Citrix**
- **Lotus Notes**
- **Microsoft Outlook**
- **Microsoft Sharepoint**
- **NetBIOS file browsing**

SAML Single Sign-on

Ivanti Connect Secure SAML 2.0 SSO Solutions

This section provides a brief overview of the Security Assertion Markup Language (SAML) standard produced and approved by the Organization for the Advancement of Structured Information Standards (OASIS).

Understanding SAML 2.0

This topic provides a reference to the Security Assertion Markup Language (SAML) standard and an overview of SAML 2.0 use cases.

About SAML

SAML is an XML-based framework for communicating user authentication, entitlement, and attribute information. The standard defines the XML-based assertions, protocols, bindings, and profiles used in communication between SAML entities. SAML is used primarily to implement Web browser single sign-on (SSO). SAML enables businesses to leverage an identity-based security system like Ivanti Connect Secure to enforce secure access to web sites and other resources without prompting the user with more than one authentication challenge.

For complete details on the SAML standard, see the OASIS web site:

http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security

SAML Use Cases

This section provides a brief summary of the primary SAML use cases.

SAML SSO

SAML is primarily used to enable Web browser single sign-on (SSO). Pulse client R5.0 or greater (Win and Mac) also supports SAML SSO. The user experience objective for SSO is to allow a user to authenticate once and gain access to separately secured systems without resubmitting credentials. The security objective is to ensure the authentication requirements are met at each security checkpoint.

In an SSO transaction, the SAML services implemented at each secured system exchange requests and assertions to determine whether to allow access. The SAML assertions used in SSO transactions include authentication statements and attribute statements.

SAML ACL

SAML can also be used to enforce access control list (ACL) permissions. In an ACL transaction, the SAML services implemented for each secured system exchange assertions to determine whether a user can access the resource. The SAML assertions used in ACL transactions include authorization requests and authorization decision statements.

SAML 2.0 Supported Features Reference

This topic provides an overview of Ivanti Connect Secure support for Security Assertion Markup Language (SAML) single sign-on (SSO). It includes the following information related to SAML 2.0 support:

- [Supported SAML SSO Deployment Modes](#)
- [Supported SAML SSO Profiles](#)
- [FIPS Support Notes](#)

Supported SAML SSO Deployment Modes

In a SAML deployment, a SAML service provider is a secured resource (an application, web site, or service) that is configured to request authentication from a SAML identity provider. The SAML identity provider responds with assertions regarding the identity, attributes, and entitlements (according to your configuration). The exchange enforces security and enables the SSO user experience.

The system can act as a SAML service provider, a SAML identity provider, or both. The following sections provide illustrations:

- [Ivanti Connect Secure as a SAML Service Provider](#)
- [Ivanti Connect Secure As a SAML Identity Provider \(Gateway Mode\)](#)
- [Ivanti Connect Secure as a SAML Identity Provider \(Peer Mode\)](#)

Ivanti Connect Secure as a SAML Service Provider


If you are working with a partner that has implemented a SAML identity provider, you can deploy the system as a SAML service provider to interoperate with it, thereby enabling SSO for users who should have access to resources in both networks. In this model, the user is authenticated by the SAML identity provider. The system uses the SAML response containing the assertion to make an authentication decision.

The choices the identity provider makes to implement SAML determine the deployment choices, for example whether to use SAML 2.0 or SAML 1.1, whether to reference a published metadata configuration file, and whether to use a POST or artifact profile. When you deploy the system as a SAML service provider, you create a SAML authentication server configuration that references the partner SAML identity provider, and a set of access management framework objects (realm, role mapping rules, and sign-in policy) that reference the SAML authentication server.

When you configure the SAML service provider, some particular settings are necessary to support either identity-provider-initiated or service-provider-initiated SSO. The documentation for the configuration steps makes note of these settings. Regardless, you configure the SAML service provider to support both identity-provider-initiated and service-provider-initiated SSO.

The following figure illustrates the flow of network communication in a service-provider-initiated SSO scenario with a Web browser client.

The following figure depicts the Ivanti Connect Secure as a SAML Service Provider in a Service-Provider-Initiated SSO Scenario:



1 - The user clicks a link to access a resource.

2a - The service provider sends an HTTP redirect status code (HTTP 302) to the user. The SAML request and all other SAML details are sent as URL parameters in the URL Location header.

2b - The user sends an HTTP GET request to the identity provider. The SAML request and all other SAML details are sent as URL parameters.

If the user already has a session with the identity provider, steps 2.1 and 2.2 are skipped.

2.1 - If the user does not have a session, the identity provider sends an authentication challenge to the user.

2.2 - The user enters sign-in credentials.


3a - The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.

3b - The user sends the form to the service provider.

4 - The external resource is delivered to the user's browser.

The following figure illustrates the flow of network communication in an identity-provider-initiated SSO scenario with a Web browser client.

The following figure depicts the Ivanti Connect Secure as a SAML Service Provider in an Identity-Provider-Initiated SSO Scenario:



1 - The user authenticates to the identity provider.

2 - The identity provider returns a portal page with links to external resources.

3 - The user clicks a link for an external resource.


4a - The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.

4b - The user sends the form to the service provider.

5 - The external resource is delivered to the user's browser.

The following figure illustrates the flow of network communication when a user clicks a Pulse client connection.

The figure depicts the Ivanti Connect Secure as a SAML Service Provider in a Pulse-Client-Initiated Connection:



1 - The user clicks the Pulse client connection. The Pulse client and system exchange IF-T/TLS messages. The Pulse client learns that authentication is a SAML exchange, and Ivanti launches its embedded client Web browser.

2a - The service provider sends an HTTP redirect status code (HTTP 302) to the user. The SAML request and all other SAML details are sent as URL parameters in the URL Location header.

2b - The user sends an HTTP GET request to the identity provider. The SAML request and all other SAML details are sent as URL parameters.

If the user already has a session with the identity provider, steps 2.1 and 2.2 are skipped.

2.1 - If the user does not have a session, the identity provider sends an authentication challenge to the user.

2.2 - The user enters sign-in credentials.

3a - The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.

3b - The user sends the form to the service provider.

4 - The setup client is run on the endpoint, and the Pulse client and system set up an SSL VPN tunnel.

Ivanti Connect Secure As a SAML Identity Provider (Gateway Mode)


When you deploy the system in front of enterprise resources that support SAML and have been configured as a SAML service provider, the system acts as a gateway for user access to the secured resource, just as it does with its other resource policies. In the SAML exchange, the system acts as a SAML identity provider. When deployed as a gateway, the SAML SSO communication is always identity-provider-initiated. The system maintains the session and uses its rewriting or pass-through proxy features to render data to the user.

In a gateway mode deployment, you configure the system as a SAML identity provider to correspond with the SAML service provider, and you create a SAML SSO resource policy configuration to determine the users and resources to which the SAML SSO experience applies. The SAML SSO resource policy supports two types of behavior that are possible with the HTTP responses sent by SAML service providers:

- The SAML service provider sends HTTP responses that can be handled by HTTP cookies and therefore do not require user interaction. In this case, the SAML SSO resource policy can be configured to use cookies to handle the HTTP transaction.
- The SAML service provider sends HTTP responses that require user interaction. For example, the SAML service provider might send an HTTP 200 OK with an embedded form that requires action from the user, execution of JavaScript, or data to be automatically submitted on load. Or, the resource might send an HTTP 3xx redirect that requires acceptance by the user. In these cases, the SAML SSO resource policy can be configured to forward the HTTP responses through the rewriter, which rewrites the HTTP response and sends it to the end user.

The following figure illustrates the communication that occurs when the SAML SSO policy is configured to handle the SAML service provider responses using cookies.

Ivanti Connect Secure as a SAML Identity Provider (Gateway Mode) - User/Browser Action Not Required:



1 - User requests a SAML protected resource.

2 - The system executes the SAML SSO policy and the identity provider sends an HTTP request containing the SAML assertion to the SAML service provider.

3 - The SAML service provider sends an HTTP response. The SAML SSO process extracts the cookies from the response and stores them in the cookie cache.


4 - The system rewriter process sends the request for the resource (sending the cookies received in step 3).

5 - The SAML service provider sends the resource.

6 - The system rewrites the resource and sends it to the user.

The following figure illustrates the communication that occurs when the SAML SSO policy is configured to rewrite the SAML service provider responses and send them to the user/browser for action.

The following figure depicts the Ivanti Connect Secure as a SAML Identity Provider (Gateway Mode) - User/Browser Action Required:



1 - User requests a SAML protected resource.

2 - The system executes the SAML SSO policy and the system identity provider sends an HTTP request containing the SAML assertion to the SAML service provider.

3 - The SAML service provider sends an HTTP response. The system SAML SSO process forwards the entire response to the rewriter.

3.1 - The rewriter rewrites the response and sends it to the user.

3.2 - The user/browser completes any action required and sends a response (an HTTP GET/POST request).

4 - The rewriter processes it as any other HTTP web request and forwards to the SAML service provider.

5 - The SAML service provider sends the resource.

6 - The system rewrites the resource and sends it to the user. Steps 5 and 6 can involve many transactions related to Web browsing or use of the resource.

Ivanti Connect Secure as a SAML Identity Provider (Peer Mode)

When deployed to support access to external resources (for example, public cloud resources), the system does not have to be a gateway to user access. The user can access the external resource directly, and the traffic does not flow through the device. In a peer mode deployment, you configure the system as a SAML identity provider to correspond with the external SAML service provider, and you create a SAML External Apps SSO resource policy configuration to determine the users and resources to which the SAML SSO experience applies.


In a service-provider-initiated SSO scenario, the user requests a resource protected by the SAML service provider. The SAML service provider redirects the user to the sign-in page. The access management framework processes the authentication request, performs host checking rules and role mapping rules. If authentication is successful, the system redirects the user to the protected resource.

In an identity-provider-initiated SSO scenario, the user first creates a session. The access management framework processes are run when the user signs in. The SAML External Apps SSO policy is enforced when the user browses to the SAML protected external application.

When you configure the SAML identity provider, some settings are necessary to support either identity-provider-initiated or service-provider-initiated SSO. The documentation for the configuration steps makes note of these settings. Regardless, you configure the SAML identity provider to support both identity-provider-initiated and service-provider-initiated SSO.

The following figure illustrates the flow of network communication in a service-provider-initiated SSO scenario.

Ivanti Connect Secure as a SAML Identity Provider (Peer Mode) in a Service-Provider-Initiated SSO Scenario:



1 - The user clicks a link to access a resource.

2a - The service provider sends an HTTP redirect status code (HTTP 302) to the user. The SAML request and all other SAML details are sent as URL parameters in the URL Location header.

2b - The user sends an HTTP GET request to the identity provider. The SAML request and all other SAML details are sent as URL parameters.

If the user already has a session with the identity provider, steps 2.1 and 2.2 are skipped.

2.1 - If the user does not have a session, the identity provider sends an authentication challenge to the user.

2.2 - The user enters sign-in credentials.


3a - The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.

3b - The user sends the form to the service provider.

4 - The external resource is delivered to the user's browser.

The following figure illustrates the flow of network communication in an identity-provider-initiated SSO scenario.

Ivanti Connect Secure as a SAML Identity Provider (Peer Mode) in an Identity-Provider-Initiated SSO Scenario:



1 - The user authenticates to the identity provider.

2 - The identity provider returns a portal page with links to external resources.

3 - The user clicks a link for an external resource.

4a - The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.

4b - The user sends the form to the service provider.

5 - The external resource is delivered to the user's browser.

Supported SAML SSO Profiles

The following table summarizes support for SAML 2.0 deployment profiles.

Profile	Message Flows	Binding	Service Provider	Identity Provider (Gateway)	Identity Provider (Peer)
Web browser SSO	<AuthnRequest> from service provider to identity provider	HTTP redirect	Supported	-	Supported
		HTTP POST	Not supported	-	Supported
		HTTP artifact	Not supported	-	Not supported
	Identity provider <Response> to service provider	HTTP POST	Supported	Supported	Supported
		HTTP artifact	Supported	Supported	Supported
	Authentication context classes sent in <RequestedAuthnContext>	HTTP POST	All authentication context classes	-	Password-Protected, TimeSyncToken, and TLSClient only
HTTP artifact		All authentication context classes	-	Password-Protected, TimeSyncToken, and TLSClient only	
Basic attribute profile	Simple <Attribute> statements (name-value pairs) as part of assertions	HTTP POST	Consumes and stores Attribute statements	Sends Attribute statements	Sends Attribute statements
		HTTP artifact			

Profile	Message Flows	Binding	Service Provider	Identity Provider (Gateway)	Identity Provider (Peer)
Assertion query / request	Artifact resolution <ArtifactResolve> <ArtifactResponse>	SOAP	Supported	Supported	Supported
Single logout	Logout request	HTTP redirect	Supported	Not supported	Not supported
		HTTP POST	Not supported	Not supported	Not supported
		HTTP artifact	Not supported	Not supported	Not supported
		SOAP	Not supported	Not supported	Not supported
	Logout response	HTTP redirect	Supported	Not supported	Not supported
		HTTP POST	Not supported	Not supported	Not supported
		HTTP artifact	Not supported	Not supported	Not supported
		SOAP	Not supported	Not supported	Not supported



Ivanti Connect Secure does not act as an Attribute Authority.

FIPS Support Notes

Historically, in FIPS deployments, private keys were managed in a way that can be problematic for SAML functionality that depends on access to the private key. The following table summarizes our support for SAML for FIPS deployments.

The following table describes the SAML Support for FIPS Deployments

	SAML 2.0		SAML 1.1	
Release	Service Provider	Identity Provider	Consumer	Producer
8.0 and above	Device certificate signing is supported; however, the ECDSA certificates is not supported. There are no other limitations.	Device certificate signing is supported; however, the ECDSA certificates is not supported. There are no other limitations.	No limitations	Artifact profile only

SAML 2.0 Configuration Tasks

This section includes the tasks you perform to enable and configure SAML services.

Configuring System-Wide SAML Settings

This section describes tasks related to configuring system-wide SAML settings.

Configuring Global SAML Settings

The system-wide SAML settings impact all SAML service provider and identity provider instances.

To configure global SAML settings:

1. Select **System > Configuration > SAML**.
2. Click the Settings button to display the configuration page.
3. Complete the settings described in the following table.
4. Click **Save Changes**.

The following table lists SAML Global Configuration Guidelines

Settings	Guidelines
Timeout value for metadata fetch request	Specify the number of seconds after which a download request is abandoned. If the peer SAML entity publishes its metadata at a remote location, the system downloads the metadata file from the specified location.
Validity of uploaded/downloaded metadata file	Specify the maximum duration for which the system considers the metadata file of the peer SAML entity to be valid. If the metadata file provided by the peer SAML entity contains validity information, the lower value takes precedence.
Host FQDN for SAML	<p>Specify the fully qualified domain name for the Ivanti Connect Secure host. The value you specify here is used in the SAML entity ID and the URLs for SAML services, including:</p> <ul style="list-style-type: none">Entity ID for SAML service provider and SAML identity provider instances. The SAML entity ID is the URL where the system publishes its SAML metadata file.Single sign-on service URLSingle logout service URLAssertion consumer service URLArtifact resolution service URL <p>BEST PRACTICE: The system uses HTTPS for these services. Therefore, we recommend that you assign a valid certificate to the interface that has the IP address to which this FQDN resolves so that users do not see invalid certificate warnings.</p>

Settings	Guidelines
Alternate Host FQDN for SAML	<p>Optional.</p> <p>If you have enabled the Reuse Existing NC (Ivanti) Session on the SAML Identity Provider Sign-In page, specify the fully qualified domain name used to generate the SSO Service URL.</p> <p>Set up your DNS service to ensure that the alternate hostname resolves to a different IP address when a session is established and when not established. We recommend the following DNS behavior:</p> <p>If the NC (Ivanti) session is not established, the IP address of the alternate hostname should resolve to the public IP address on the device external port.</p> <p>If the NC (Ivanti) session is established, the IP address of the alternate hostname should resolve to the private IP address on the device internal port.</p> <p>BEST PRACTICE: The system uses HTTPS for this service. Therefore, we recommend that you assign a valid certificate to the interface that has the IP address to which this FQDN resolves so that users do not see invalid certificate warnings.</p>
Update Entity IDs	Use this button to regenerate the SAML entity IDs of all configured service providers and identity providers. Typically, you take this action when the Host FQDN for SAML is changed.

Managing SAML Metadata Files

You use the System > Configuration > SAML pages to maintain a table of SAML metadata files for the SAML service providers and identity providers in your network. Using SAML metadata files makes configuration easier and less prone to error.

You can add the metadata files to the system by:

- Uploading a metadata file.
- Retrieving the metadata file from a well-known URL.

To add metadata files:

1. Select **System > Configuration > SAML**.
2. Click **New Metadata Provider** to display the configuration page.
3. Complete the settings described in the following table.

4. Save the configuration.

The following table lists the SAML Metadata Provider Configuration Guidelines:

Settings	Guidelines
Metadata Provider Location Configuration	Select one of the following methods: Local. Browse and locate the metadata file on your local host or file system. Remote. Enter the URL of the metadata file. Only http and https protocols are supported.
Metadata Provider Verification Configuration	
Accept Untrusted Server Certificate	If you specify a URL for the metadata provider, select this option to allow the system to download the metadata file even if the server certificate is not trusted. This is necessary only for HTTPS URLs.
Accept Unsigned Metadata	If this option is not selected, unsigned metadata is not imported. Signed metadata is imported only after signature verification.
Signing Certificate	Browse and locate the certificate that verifies the signature in the metadata file. This certificate overrides the certificate specified in the signature of the received metadata. If no certificate is uploaded here, then the certificate present in the signature of the received metadata is used. Select the Enable Certificate Status Checking option to verify the certificate before using it. Certificate verification applies both to the certificate specified here and the certificate specified in the signature in the metadata file.
Metadata Provider Filter Configuration	
Roles	Select whether the metadata file includes configuration details for a SAML service provider, identity provider, or Policy Decision Point. You may select more than one. If you select a role that is not in the metadata file, it is ignored. If none of the selected roles are present in the metadata file, the system returns an error.
Entity IDs To Import	Enter the SAML Entity IDs to import from the metadata files. Enter only one ID per line. Leave this field blank to import all IDs. This option is available only for uploading local metadata files.

The Refresh button downloads the metadata files from the remote location even if these files have not been modified. This operation applies only to remote locations; local metadata providers are ignored if selected.

To refresh a metadata file:

1. Select **System > Configuration > SAML**.
2. Select the metadata file to refresh and click **Refresh**.
3. To delete a metadata file:
4. Select **System > Configuration > SAML**.
5. Select the metadata file to delete and click **Delete**.

Configuring Ivanti Connect Secure as a SAML 2.0 Service Provider

This topic describes how to configure the system as a SAML service provider. When the system is a SAML service provider, it relies on the SAML identity provider authentication and attribute assertions when users attempt to sign in to the device. Note that authentication is only part of the security system. The access management framework determines access to the system and protected resources.

The system supports:

- HTTP Redirect binding for sending AuthnRequests
- HTTP Redirect binding for sending/receiving SingleLogout requests/responses
- HTTP POST and HTTP Artifact bindings for receiving SAML responses
- RequestedAuthnContext context class specifications

Before you begin:

- Check to see whether the SAML identity provider uses HTTP POST or HTTP Artifact bindings for SAML assertions.
- Check to see whether the SAML identity provider has published a SAML metadata file that defines its configuration. If the SAML identity provider metadata file is available, configuration is simpler and less prone to error.
- Complete the system-wide SAML settings if you have not already done so. Select **System > Configuration > SAML > Settings**. For details, see [Configuring Global SAML Settings](#)
- Add metadata for the SAML identity provider to the metadata provider list if you have not already done so. Select **System > Configuration > SAML**. For details, see [Managing SAML Metadata Files](#).

The sign-in URL for which a session needs to be established for the system as a service provider is identified by the RelayState parameter (HTTP URL parameter for artifact and HTML form parameter for POST.) In a service provider initiated case, the system populates RelayState as an HTTP URL parameter while sending AuthnRequest. In the IdP-Initiated scenario (Ivanti Connect Secure is a service provider and there is a third-party IdP), the IdP must be configured to set the appropriate Sign-in URL of the system in the RelayState parameter of the HTML form containing the SAML response. For more information, see the SAML 2.0 specification.

To configure the system as a SAML service provider:

1. Select **Authentication > Auth. Servers**.
2. Select SAML Server from the New list and then click New Server to display the configuration page.
3. Complete the settings as described in the [Table](#).
4. Save the configuration.

After you save changes for the first time, the page is redisplayed and now has two tabs. Use the Settings tab to modify any of the settings pertaining to the SAML server configuration. Use the Users tab to monitor user sessions.

Next steps:

1. Configure the access management framework to use the SAML authentication server. Start with realm and role mapping rules. For details, see "[Creating an Authentication Realm](#)" and [Specifying Role Mapping Rules for an Authentication Realm](#)
2. Configure a sign-in policy. When using a SAML authentication server, the sign-in policy can map to a single realm only. For details, see [Defining a Sign-In Policy](#)

The following table lists the SAML Service Provider Profile:

Settings	Guidelines
Name	Specify a name to identify the server instance.
Settings	
SAML Version	Select 2.0.
SA Entity Id	This value is prepopulated. It is generated by the system, based on the value for the Host FQDN for SAML setting on the System > Configuration > SAML > Settings page.

Settings	Guidelines
Configuration Mode	<p>Select Manual or Metadata. If a metadata file or location is available from the SAML identity provider, use the metadata option to make configuration simpler and less prone to error. To upload or set the location for the published metadata file, select System > Configuration > SAML and click the New Metadata Provider button.</p>
Identity Provider Entity ID	<p>The identity provider entity ID is sent as the Issuer value in the assertion generated by the SAML identity provider.</p> <p>If you use the metadata option, this setting can be completed by selecting the identity provider entity ID from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page.</p> <p>If you complete this setting manually, specify the Issuer value in assertions generated by the SAML identity provider. Typically, you ask the SAML identity provider administrator for this setting.</p>
Identity Provider Single Sign-On Service URL	<p>The identity provider SSO service URL is a URL provisioned by the SAML identity provider. The setting is required to support service-provider-initiated SSO. If missing, the system cannot successfully redirect the user request.</p> <p>If you use the metadata option, this setting can be completed by selecting the SSO service URL from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page.</p> <p>If you complete this setting manually, ask the SAML identity provider administrator for this setting.</p>

Settings	Guidelines
User Name Template	<p>Specify how the system is to derive the username from the assertion. If the field is left blank, it uses the string received in the NameID field of the incoming assertion as the username.</p> <p>If you choose a certificate attribute with more than one value, the system uses the first matched value. For example, if you enter <certDN.OU> and the user has two values for the attribute (ou=management, ou=sales), the system uses "management". To use all values, add the SEP attribute to the variable. For example, if you enter <certDN.OUT SEP=":">, the system uses "management:sales". The attributes received in the attribute statement in the incoming assertion are saved under userAttr. These variables can also be used with angle brackets and plain text. If the username cannot be generated using the specified template, the login fails. If the NameID field of the incoming assertion is of type X509Nameformat, then the individual fields can be extracted using system variable "assertionNameDN".</p>
Allowed Clock Skew (minutes)	<p>Specify the maximum allowed difference in time between the system clock and the SAML identity provider server clock.</p> <p>NOTE: SAML is a time sensitive protocol. The time-based validity of a SAML assertion is determined by the SAML identity provider. If the SAML identity provider and SAML service provider clocks are askew, the assertion can be determined invalid, and you will receive the following error:</p> <p>"SAML Transferred failed. Please contact your system administrator. Detail: Failure: No valid assertion found in SAML response."</p> <p>We recommend you use NTP to ensure the clocks are synchronized and that you set an Allowed Clock Skew value that accommodates any expected or permissible skew.</p>

Settings	Guidelines
Support Single Logout	<p>Single logout is a mechanism provided by SAML for logging out a particular user from all the sessions created by the identity provider. Select this option if the system must receive and send a single logout request for the peer SAML identity provider.</p> <p>If you use the metadata option, the Single Logout Service URL setting can be completed by selecting the SLO service URL from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page. The system sends Single Logout requests to this URL.</p> <p>In addition, if you use the metadata option, the Single Logout Response URL setting is completed based on your selection for Single Logout Service URL. If the identity provider has left this setting empty in its metadata file, the system sends the Single Logout response to the SLO service URL.</p> <p>If you complete these settings manually, ask the SAML identity provider administrator for guidance.</p>
	<p>The Support Single Logout service for the identity provider must present a valid certificate. For example, the hostname in a single logout request URL must be the same as the Common Name of the certificate presented by the identity provider of that hostname. If an invalid certificate is presented, the single logout feature may not work as intended.</p>
SSO Method	

Settings	Guidelines
Artifact	<p>When configured to use the Artifact binding, the system contacts the Artifact Resolution Service (ARS) to fetch the assertion using SOAP protocol. If the ARS is hosted on a HTTPS URL, then the certificate presented by the ARS is verified by the system. For this verification to pass successfully, the CA of the server certificate issued to the identity provider ARS must be added to the trusted server CA on the system.</p> <p>Complete the following settings to configure SAML using the HTTP Artifact binding:</p> <p>Source ID. Enter the source ID for the identity provider ARS. Source ID is Base64-encoded, 20-byte identifier for the identity provider ARS. If left blank, this value is generated by the system.</p> <p>Source Artifact Resolution Service URL. For metadata-based configuration, this field is completed automatically from the metadata file and is not configurable. For manual configurations, enter the URL of the service to which the SP ACS is to send ArtifactResolve requests. ArtifactResolve requests are used to fetch the assertion from the artifact received by it.</p> <p>SOAP Client Authentication. Select HTTP Basic or SSL Client Certificate and complete the related settings. If you use an SSL client certificate, select a certificate from the device certificate list.</p> <p>Select Device Certificate for Signing. Select the device certificate the system uses to sign the AuthnRequest sent to the identity provider SSO service. If you do not select a certificate, the system does not sign AuthnRequest.</p> <p>Select Device Certificate for Encryption. Select the device certificate the system uses to decrypt encrypted data received in the SAML response. The public key associated with the device certificate is used by the identity provider for encryption.</p>

Settings	Guidelines
POST	<p>When configured to use the POST binding, the system uses a response signing certificate to verify the signature in the incoming response or assertion. The certificate file must be in PEM or DER format. The certificate you select should be the same certificate used by the identity provider to sign SAML responses.</p> <p>Complete the following settings to configure SAML using the HTTP POST binding:</p> <p>Response Signing Certificate. If you use the metadata-based configuration option, select a certificate from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page.</p> <p>If you configure these settings manually, browse to and upload the certificate to be used to validate the signature in the incoming response or assertion. If no certificate is specified, the certificate embedded in the response is used.</p> <p>Enable Signing Certificate status checking. Select this option to check the validity of the signing certificate before verifying the signature. This setting applies to any certificate used for signature verification. If this option is enabled, the response will be rejected if the certificate is revoked, expired, or untrusted. If this option is selected, the certificate CA must be added to the system Trusted Client CA store.</p> <p>If this option is not enabled, then the certificate is used without any checks.</p> <p>Select Device Certificate for Signing. Select the device certificate the system uses to sign the AuthnRequest sent to the identity provider SSO service. If you do not select a certificate, the system does not sign AuthnRequest.</p> <p>Select Device Certificate for Encryption. Select the device certificate the system uses to decrypt encrypted data received in the SAML response. The public key associated with the device certificate is used by the identity provider for encryption.</p>

Settings	Guidelines
Authentication Context Classes	<p>Use the Add and Remove buttons to select authentication context classes to be sent in the authentication requests to the SAML identity provider. These are included in the RequestedAuthnContext element.</p> <p>In the OASIS standard, an authentication context is defined as "the information, additional to the authentication assertion itself, that the relying party may require before it makes an entitlements decision with respect to an authentication assertion."</p> <p>This feature supports all authentication context classes specified in the SAML 2.0 OASIS Authn Context specification.</p> <p>For example, if you select X509, the system sends the following context:</p> <pre><samlp:RequestedAuthnContext> <saml:AuthnContextClassRef xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion"> urn:oasis:names:tc:SAML:2.0:ac:classes:X509</saml:AuthnContextClassRef> </samlp:RequestedAuthnContext></pre> <p>In response, the SAML IdP sends the context data along with the authentication results. The system stores the context data in the session cache, and it can be specified in user attribute role mapping rules.</p> <p>Specify a comparison attribute within the RequestedAuthnContext element. The comparison attribute specifies the relative strengths of the authentication context classes specified in the request and the authentication methods offered by a SAML IdP. The following values specified in the SAML 2.0 OASIS core specification can be selected:</p> <ul style="list-style-type: none"> exact - Requires the resulting authentication context in the authentication statement to be the exact match of at least one of the authentication contexts specified. minimum - Requires the resulting authentication context in the authentication statement to be at least as strong as one of the authentication contexts specified. maximum - Requires the resulting authentication context in the authentication statement to be stronger than any one of the authentication contexts specified. better - Requires the resulting authentication context in the authentication statement to be as strong as possible without exceeding the strength of at least one of the authentication contexts specified. <p>Select the same value that is configured on the SAML IdP. If none is specified in the SAML IdP configuration, the implicit default is exact.</p>

Settings	Guidelines
Service Provider Metadata Settings	
Metadata Validity	Enter the number of days the system metadata is valid. Valid values are 0 to 9999. 0 specifies the metadata does not expire.
Do Not Publish SA Metadata	Select this option if you do not want the system to publish the metadata at the location specified by the system Service Entity ID field.
Download Metadata	This button appears only after you have saved the authentication server configuration. Use this button to download the metadata of the current SAML service provider.
User Record Synchronization	
Enable User Record Synchronization	Allow users to retain their bookmarks and individual preferences regardless of which device they log in to.
Logical Auth Server Name	Specify the server name if you have enabled user record synchronization.

Configuring Ivanti Connect Secure as a SAML 2.0 Identity Provider

This topic describes how to configure the system as a SAML identity provider.

Configuration Overview

Implementing the system as a SAML identity provider includes the following basic steps.

1. Configure system-wide SAML settings. Select **System > Configuration > SAML > Settings**. See [Configuring Global SAML Settings](#).
2. Add SAML metadata provider files. Select **System > Configuration > SAML**. See [Managing SAML Metadata Files](#).
3. Configure Sign-In SAML metadata provider settings. See [Configuring Sign-in SAML Metadata Provider Settings](#).
4. Configure Sign-In SAML identity provider settings. See [Configuring Sign-in SAML Identity Provider Settings](#).
5. Configure peer service provider settings. See [Configuring Peer SAML Service Provider Settings](#).

6. Configure a resource policy:
 - For gateway mode deployments, configure a SAML SSO resource policy. See [Configuring a SAML SSO Resource Policy for Gateway Mode Deployments](#)
 - For peer mode deployments, configure a SAML SSO external applications policy. See [Configuring a SAML External Applications SSO Policy](#)

Configuring Sign-in SAML Metadata Provider Settings

Sign-in SAML metadata provider settings determine how the system identity provider metadata is published.

To configure the identity provider metadata publication settings:

1. Select **Authentication > Signing In > Sign-In SAML > Metadata Provider** to display the configuration page.
2. Complete the settings described in the following table.
3. Click **Save Metadata Provider** to save your changes.

The following table lists the Sign-in SAML Identity Provider Metadata Provider Configuration Guidelines:

Settings	Guidelines
Entity ID	This value is prepopulated. It is generated by the system, based on the value for the Host FQDN for SAML setting on the System > Configuration > SAML > Settings page.
Metadata Validity	Specify the maximum duration for which a peer SAML entity can cache the system SAML metadata file. Valid values are 1 to 9999. The default is 365 days.
Do Not Publish SA Metadata	Select this option if you do not want the system to publish the metadata at the location specified by the system Entity ID field. You can use this option to toggle off publication without deleting your settings.
Download Metadata	Use this button to download the system SAML identity provider metadata.

Configuring Sign-in SAML Identity Provider Settings

The settings defined in this procedure are the default settings for the system SAML identity provider communication with all SAML service providers. If necessary, you can use the peer service provider configuration to override these settings for particular service providers.

To configure sign-in SAML identity provider settings:

1. Select **Authentication > Signing In > Sign-In SAML > Identity Provider** to display the configuration page.
2. Complete the settings described in the following table.
3. Save the configuration.

The following table lists the Sign-in SAML Identity Provider Configuration Guidelines:

Settings	Guidelines
Basic Identity Provider (IdP) Configuration (Published in Metadata)	
Protocol Binding to use for SAML Response	Select POST, Artifact, or both, depending on your total requirements.
Signing Certificate	Select the certificate used to sign the SAML messages sent by the system. The certificates listed here are configured on the System > Configuration > Certificate > Device Certificates page.
Decryption Certificate	Select the certificate used to decrypt the SAML messages sent by peer service providers. The public key associated with this certificate is used by the peer service provider to encrypt SAML messages exchanged with this identity provider. The decryption certificate must be configured if the peer service provider encrypts the SAML messages sent to the system. The certificates listed here are configured on the System > Configuration > Certificate > Device Certificates page.

Settings	Guidelines
Other Configurations	<p>Reuse Existing NC (Ivanti) Session. This feature applies to a service-provider-initiated SSO scenario - that is, when a user clicks a link to log into the service provider site. The service provider redirects the user to the identity provider SSO Service URL.</p> <p>If this option is selected, a user with an active NC/Ivanti session is not prompted to authenticate. The system uses information from the existing session to form the SAML response.</p> <p>Accept unsigned AuthnRequest. In a service-provider-initiated SSO scenario, the SP sends an AuthnRequest to the identity provider. This AuthnRequest could be either signed or unsigned. If this option is unchecked, the system rejects unsigned AuthnRequests. Note that the system also rejects signed AuthnRequests if signature verification fails.</p>
Service-Provider-Related IdP Configuration	
Relay State	SAML RelayState attribute sent to the service provider in an identity-provider-initiated SSO scenario. If left blank, the RelayState value is the URL identifier of the resource being accessed.
Session Lifetime	<p>Suggest a maximum duration of the session at the service provider created as a result of the SAML SSO. Select one of the following options:</p> <p>None. The identity provider does not suggest a session duration.</p> <p>Role Based. Suggest the value of the session lifetime configured for the user role.</p> <p>Customized. If you select this option, the user interface displays a text box in which you specify a maximum in minutes.</p>
Sign-In Policy	<p>Select the sign-in URL to which the user is redirected in a service-provider-initiated scenario. The list is populated by the sign-in pages configured on the Authentication > Signing In > Sign-in Policies page.</p> <p>The user is not redirected if he or she already has a session with the system and had authenticated through this sign-in policy.</p>

Settings	Guidelines
Force Authentication Behavior	<p>In an service-provider-initiated scenario, the service provider sends an AuthnRequest to the identity provider. If the service provider AuthnRequest includes the ForceAuthn attribute set to true and the user has a valid session, this setting determines how the identity provider responds. Select one of the following options:</p> <p>Reject AuthnRequest. Do not honor the SAML SSO request.</p> <p>Re-Authenticate User. Invalidate the user session and prompt for reauthentication.</p> <p>This setting prevails over the Ivanti session reuse setting.</p>
User Identity	
Subject Name Format	<p>Format of the NameIdentifier field in the generated assertion. Select one of the following options:</p> <p>DN. Username in the format of DN (distinguished name).</p> <p>Email address. Username in the format of an e-mail address.</p> <p>Windows. Username in the format of a Windows domain qualified username.</p> <p>Other. Username in an unspecified format.</p>
Subject Name	<p>Template for generating the username that is sent as the value of the NameIdentifier field in the assertion.</p> <p>You may use any combination of available system or custom variables contained in angle brackets and plain text.</p>
Web Service Authentication	<p>These settings apply when the HTTP Artifact binding is used.</p>
Authentication Type	<p>Method used to authenticate the service provider assertion consumer service to the identity provider on the system. Select one of the following options:</p> <p>None. Do not authenticate the assertion consumer service.</p> <p>Username/Password. If you select this option, use the controls to specify username and password settings.</p> <p>Certificate. For certificate-based authentication, the Client CA of the service provider should be present in the system Trusted Client CA list (located on the System > Configuration > Certificates > Trusted Client CAs page).</p>
Artifact Configuration	<p>These settings apply when the HTTP Artifact binding is used.</p>

Settings	Guidelines
Source ID	This is the Base64-encoded, 20-byte identifier of the Artifact Resolution Service on the identity provider.
Enable Artifact Response Signing and Encryption	If checked, the identity provider signs and encrypts the artifact response.
Attribute Statement Configuration	Attributes to be sent in SAML Attribute statements can be specified manually as name-value pairs, or you can configure an option to fetch name-value pairs from an LDAP server (or you can specify both manual entries and LDAP entries).
Attribute Name	An ASCII string.
Friendly Name	A more readable friendly name for the attribute. This is optional (an option included in the SAML standard).
Attribute Value	<p>The attribute value can be specified as a hard-coded string, a custom variable, or a user attribute variable. System conventions for specifying user and custom tokens and variables apply.</p> <p>The value can be a combination of a string and a user or custom variable. For example: Email::<customVar.email>. The value can also be a combination of user and custom variables and hardcoded text. For example: mydata= <USER> <REALM> <customVar.email>.</p>

Settings	Guidelines
Value Type	<p>Select Single-Valued or Multivalued.</p> <p>A single-valued attribute can be a combination of a string and a user or custom variable.</p> <p>If there are multiple single-valued attributes configured with the same attribute names, they are combined and sent as a multivalued attribute.</p> <p>Select Multivalued if you want every individual token defined in the Attribute Value column to be sent as a separate AttributeValue. For example:</p> <pre><element ref="saml:AttributeValue" minOccurs="0" maxOccurs="unbounded"/></pre> <p>If the Attribute value is given as <USER>mars<REALM>pulsesecure<ROLE> and the value type is marked as Multivalued, then the values sent as part of attribute statement are sent as follows:</p> <p>Username Realmname Role</p> <p>Note that only the tokens ['<>'] will be considered when processing a Multivalued attribute marked. The remaining data (for example mars, jupiter) is discarded.</p> <p>Specifying the token <ROLE> will send only one role. To send all roles, specify the Attribute value with the syntax <ROLE SEP=",">. If you specify <ROLE SEP=","> as a single-valued attribute, it is sent as a single string with "," separated roles. If you specify <ROLE SEP=","> as a multi-valued attribute, each role is sent in a separate <AttributeValue> element.</p> <p>Encryption is set at the assertion level. You cannot encrypt individual attributes.</p>

Settings	Guidelines
Directory Server	<p>To fetch attribute name-value pairs from an LDAP server, complete the following settings:</p> <p>Directory Server - Select the LDAP server from the list. You must add the LDAP server to the Authentication > Auth. Servers list before it can be selected.</p> <p>Username for lookup - Enter a username template for LDAP lookup. The default is the variable <USERNAME>. The <USERNAME> variable stands for the login credential the user entered when logging in. The value can contain contextual characters as well as variables for substitution.</p> <p>Attribute Name - Type an LDAP attribute name, such as cn. The attribute name is fetched from the LDAP server and sent as SAML Attribute statements as part of a SAML assertion.</p> <p>Friendly Name - A more readable friendly name for the attribute. This is optional (an option included in the SAML standard).</p> <p>With the LDAP option, the SAML IdP sends attributes in the form configured on the backend LDAP server. If the LDAP server returns an attribute value in multi-valued form, then the SAML attribute statement will also be in multi-valued form.</p>

Configuring Peer SAML Service Provider Settings

The peer service provider list defines the set of service providers configured to communicate with the system SAML identity provider. When you add a peer service provider to the list, you can customize the SAML identity provider settings used to communicate with the individual service provider. If the service provider provides a SAML metadata file, you can use it to simplify configuration, or you can complete more detailed manual steps. If available, we recommend you use metadata so that configuration is simpler and less prone to error.

To configure peer SAML service provider settings:

1. Select **Authentication > Signing In > Sign-In SAML > Identity Provider**.
2. Under **Peer Service Provider Configuration**, create a list of service providers that are SAML peers to the system **SAML identity provider**. To add a service provider to the list, click **Add SP** to display the configuration page.
3. Complete the settings described in the following table.
4. Save the configuration.

The following table lists the Peer Service Provider Configuration Guidelines:

Settings	Guidelines
Configuration Mode	Select Manual or Metadata.
Service Provider Configuration - Metadata	
Entity Id	If you use metadata, select the SAML entity ID of the service provider. This list contains all the service providers specified in all the metadata files added to the System > Configuration > SAML page.
Select certificates manually	When you use the metadata configuration, the system SAML identity provider iterates through all the signature verification certificates specified when verifying the incoming SAML messages coming from the service provider. Similarly, when encrypting the SAML messages going out, the system SAML identity provider encrypts the messages with the first valid encryption certificate encountered in the metadata. Select this option to override this default behavior and select certificates manually.
Signature Verification Certificate	If you select the Select certificates manually option, select the certificate to be used by the identity provider to verify the signature of incoming SAML messages.
Encryption Certificate	If you select the Select certificates manually option, select the certificate to be used if the assertions sent by the identity provider must be encrypted.
Service Provider Configuration - Manual	
Entity Id	If you are completing a manual configuration, ask the SAML service provider administrator for this setting.
Assertion Consumer Service URL	SAML service provider URL that receives the assertion or artifact sent by the identity provider.
Protocol Binding supported by the Assertion Consumer Service at the SP	Select POST, Artifact, or both. This setting must be consistent with the SAML identity provider configuration.

Settings	Guidelines
Default Binding	<p>If both POST and Artifact bindings are supported, which is the default?</p> <p>Post</p> <p>Artifact</p> <p>This setting must be consistent with the SAML identity provider configuration.</p>
Signature Verification Certificate	<p>Upload the certificate to be used by the identity provider to verify the signature of incoming SAML messages. If no certificate is specified, the certificate embedded in the incoming SAML message is used for signature verification.</p>
Encryption Certificate	<p>Upload the certificate to be used if the assertions sent by the identity provider must be encrypted. If not certificate is specified, the assertions sent by the identity provider are not encrypted.</p>
Certificate Attribute Configuration for Artifact Resolution Service	<p>Optional. Specify attributes that must be present in the certificate presented to the Artifact Resolution Service (ARS) at the identity provider by the service provider assertion consumer service.</p> <p>This option appears only if the SAML service provider supports the HTTP Artifact binding, the system SAML identity provider has been configured to support the HTTP Artifact binding, and the Web service authentication type specified for the service provider is Certificate.</p> <p>Certificate Status Checking Configuration</p>
Enable signature verification certificate status checking	<p>Select this option to enable revocation checks for the signing certificate.</p> <p>Uses the configuration on the System > Configuration > Certificates > Trusted Client CAs page.</p>
Enable encryption certificate status checking	<p>Select this option to enable revocation checks for the encryption certificate.</p> <p>Uses the configuration on the System > Configuration > Certificates > Trusted Client CAs page.</p>
Customize identity provider Behavior	
Override Default Configuration	<p>Select this option to set custom behavior of the system SAML identity provider for this SP instance. If you select this option, the user interface displays the additional options listed next.</p>
Reuse Existing NC (Ivanti) Session	<p>This option cannot be enabled here if it is not selected for the sign-in SAML identity provider default settings.</p>

Settings	Guidelines
Accept unsigned AuthnRequest	Individual service providers can choose to accept unsigned AuthnRequest.
Relay State	SAML RelayState attribute sent to the service provider in an identity-provider-initiated SSO scenario. If left blank, the RelayState value is the URL identifier of the resource being accessed.
Session Lifetime	Suggest a maximum duration of the session at the service provider created as a result of the SAML SSO. Select one of the following options: None. The identity provider does not suggest a session duration. Role Based. Suggest the value of the session lifetime configured for the user role. Customized. If you select this option, the user interface displays a text box in which you specify a maximum in minutes.
Sign-In Policy	Select the Sign-In URL to which the user is redirected in a service-provider-initiated scenario. The list is populated by the sign-in pages configured in Authentication > Signing In > Sign-in Policies. The user is not redirected if he or she already has an active session and had authenticated through this sign-in policy.
Force Authentication Behavior	In an service-provider-initiated scenario, the service provider sends an AuthnRequest to the identity provider. If the service provider AuthnRequest includes the ForceAuthn attribute set to true and the user has a valid session, this setting determines how the identity provider responds. Select one of the following options: Reject AuthnRequest. Do not honor the SAML SSO request. Re-Authenticate User. Invalidate the user session and prompt for reauthentication. This setting prevails over the Ivanti session reuse setting.
User Identity	
Subject Name Format	Format of NameIdentifier field in generated Assertion. Select one of the following options: DN. Username in the format of DN (distinguished name). Email address. Username in the format of an e-mail address. Windows. Username in the format of a Windows domain qualified username. Other. Username in an unspecified format.

Settings	Guidelines
Subject Name	<p>Template for generating the username that is sent as the value of the NameIdentifier field in the assertion.</p> <p>You may use any combination of available system or custom variables contained in angle brackets and plain text.</p>
Web Service Authentication	<p>These settings apply when the HTTP Artifact binding is used.</p>
Authentication Type	<p>Method used to authenticate the service provider assertion consumer service to the identity provider on the system. Select one of the following options:</p> <p>None. Do not authenticate the assertion consumer service.</p> <p>Username/Password. Use the controls to specify username and password settings.</p> <p>Certificate. For certificate-based authentication, the client CA of the service providers should be present in the system trusted client CA list (located on the System > Configuration > Certificates > Trusted Client CAs page).</p>
Artifact Configuration	<p>These settings apply when the HTTP Artifact binding is used.</p>
Source ID	<p>This is the Base64-encoded, 20-byte identifier of the Artifact Resolution Service on the identity provider.</p>
Enable Artifact Response Signing and Encryption	<p>If checked, the identity provider signs and encrypts the Artifact response.</p>
Attribute Statement Configuration	

Settings	Guidelines
Send Attribute Statements	<p>Select this option if the SAML SP requires additional attributes to be sent with SAML assertions.</p> <p>If you enable attribute statements, select one of the following configuration options:</p> <p>Use IdP Defined Attributes-Send attributes based on the default settings for the system SAML identity provider communication with all SAML service providers.</p> <p>Customize IdP Defined Attributes-Selectively configure the attributes that are sent for this particular peer SAML SP. Attributes to be sent in SAML Attribute statements can be specified manually as name-value pairs, or you can configure an option to fetch name-value pairs from an LDAP server (or you can specify both manual entries and LDAP entries). If you select this option, configure the settings described next.</p>
Attribute Name	An ASCII string.
Friendly Name	A more readable friendly name for the attribute. This is optional (an option included in the SAML standard).
Attribute Value	<p>The attribute value can be specified as a hard-coded string, a custom variable, or a user attribute variable. System conventions for specifying user and custom tokens and variables apply.</p> <p>The value can be a combination of a string and a user or custom variable. For example: Email::<customVar.email>. The value can also be a combination of user and custom variables and hardcoded text. For example: mydata=<USER><REALM><customVar.email>.</p>

Settings	Guidelines
Value Type	<p>Select Single-Valued or Multivalued.</p> <p>A single-valued attribute can be a combination of a string and a user or custom variable.</p> <p>If there are multiple single-valued attributes configured with the same attribute names, they are combined and sent as a multivalued attribute.</p> <p>Select Multivalued if you want every individual token defined in the Attribute Value column to be sent as a separate AttributeValue. For example:</p> <pre><element ref="saml:AttributeValue" minOccurs="0" maxOccurs="unbounded"/></pre> <p>If the Attribute value is given as <USER>mars<REALM>pulsesecure<ROLE> and the value type is marked as multivalued, then the values sent as part of attribute statement are sent as follows:</p> <ul style="list-style-type: none"> Username Realmname Role <p>Note that only the tokens ['<>'] will be considered when processing a multivalued attribute marked. The remaining data (for example mars, jupiter) is discarded.</p> <p>Specifying the token <ROLE> will send only one role. To send all roles, specify the Attribute value with the syntax <ROLE SEP=",">. If you specify <ROLE SEP=","> as a single-valued attribute, it is sent as a single string with "," separated roles. If you specify <ROLE SEP=","> as a multivalued attribute, each role is sent in a separate <AttributeValue> element.</p> <p>Encryption is set at the assertion level. You cannot encrypt individual attributes.</p>

Settings	Guidelines
Directory Server	<p>To fetch attribute name-value pairs from an LDAP server, complete the following settings:</p> <p>Directory Server-Select the LDAP server from the list. You must add the LDAP server to the Authentication > Auth. Servers list before it can be selected.</p> <p>Username for lookup-Enter a username template for LDAP lookup. The default is the variable <USERNAME>. The <USERNAME> variable stands for the login credential the user entered when logging in. The value can contain contextual characters as well as variables for substitution.</p> <p>Attribute Name-Type an LDAP attribute name, such as cn. The attribute name is fetched from the LDAP server and sent as SAML Attribute statements as part of a SAML assertion.</p> <p>Friendly Name-A more readable friendly name for the attribute. This is optional (an option included in the SAML standard).</p> <p>With the LDAP option, the SAML IdP sends attributes in the form configured on the backend LDAP server. If the LDAP server returns an attribute value in multivalued form, then the SAML attribute statement will also be in multivalued form.</p>

Configuring a SAML SSO Resource Policy for Gateway Mode Deployments

When deployed as a gateway in front of enterprise resources, the SAML SSO policy acts like other resource policies. When deployed as a gateway, the SAML SSO communication can be configured as either identity-provider-initiated or service-provider-initiated. The system maintains the session and uses its rewriting or pass-through proxy features to render data to the user. You use a SAML SSO resource policy when the protected resource supports SAML SSO and has been configured as a SAML service provider.

To configure a SAML SSO resource policy:

1. Select **Users > Resource Policies > Web**.
2. Use the tabs to display the **SSO > SAML page**.
3. If your administrator view is not configured to show SAML policies, click the **Customize** button in the upper-right corner of the page and select the **SSO** and **SAML** check boxes.
4. Click **New Policy** to display the configuration page.
5. Complete the settings described in the following table.

6. Save the configuration.

The following table lists the SAML SSO Resource Policy Configuration Guidelines:


Settings	Guidelines
Name	Type a name for the policy.
Description	Type a description that would be meaningful to other administrators.
Resources	Specify the fully qualified domain name for the resources for which this policy applies. These are the resources protected at the SAML service provider.
Roles	<p>Select one of the following options:</p> <p>Policy applies to ALL roles. To apply this policy to all users</p> <p>Policy applies to SELECTED roles. To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.</p> <p>Policy applies to all roles OTHER THAN those selected below. To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.</p>
Action	<p>Select one of the following options:</p> <p>Use the SAML SSO defined below. Typically, this is the setting you use for a SAML SSO resource policy. The system SAML identity provider/SAML service provider makes the SSO request when a user tries to access a SAML resource specified in the Resources list.</p> <p>Do NOT use SAML. The system does not perform an SSO request. Use this if there is a problem with the SAML service provider and you want to allow access.</p> <p>Use Detailed Rules. Use this option to configure advanced rules.</p>
SAML SSO Details	
SAML Version	Select 2.0.
SAML SSO Type	<p>From 9.1R2 release, an administrator has the option to choose between IdP (ICS) or SP to initiate SAML single sign on</p> <p>Select the required SAML SSO Type:</p> <p>IdP-Initiated: ICS (configured as Identity Provider) initiated SAML SSO.</p> <p>SP-Initiated: Service Provider initiated SAML SSO in Rewriter mode.</p>

Settings	Guidelines
Service Provider Entity ID	Select the service provider entity ID. The service provider entity IDs listed here are configured on the Authentication > Signing In > Sign-in SAML > Identity Provider > Peer Service Provider pages.
Cookie Domain	Enter a comma-separated list of domains to which the system sends the SSO cookie.
Rewrite Response from SP	Select this option if the SAML service provider generates HTTP responses that require user/browser action, such as submission of a form, JavaScript execution, redirection to a different location, and other similar behavior. If you select this option, the system rewrites the HTTP responses sent by the SAML service provider and sends them to the user.

Configuring Service Provider Initiated SAML SSO

From 9.1R2 release, Ivanti supports SP-initiated SAML SSO when ICS is configured as IdP in gateway mode. ICS uses the existing user session in generating SAML assertion for the user for SSO.

In SP-Initiated SSO, the sequence is as follows:



1. A user logs into ICS and clicks bookmark (SP-Initiated SAML SSO resource).
2. ICS sends the request to SP.
3. SP responds with SAML AuthnRequest to ICS as the user is not authenticated to SP.
4. ICS posts SAML assertion to SP.
5. SP sends the resource to ICS.
6. ICS rewrites the resource and provides access to the user.

To configure a SAML SSO resource policy:

1. Select **Users > Resource Policies > Web**.
2. Use the tabs to display the **SSO > SAML** page.

If your administrator view is not configured to show SAML policies, click the **Customize** button in the upper-right corner of the page and select the **SSO** and **SAML** check boxes.

3. Click **New Policy** to display the configuration page.
4. In the SAML SSO Details section, select **SAML SSO Type** as **SP-Initiated**.

5. Complete other settings described in the following table.
6. Save the configuration.

Configuring a SAML External Applications SSO Policy

When deployed to support access to external resources (for example, public cloud resources), the system does not have to be a gateway to user access. The user can access the external resource directly, and the traffic does not flow through the device. To enable SAML SSO in these deployments, you configure the system as a SAML identity provider to correspond with the external SAML service provider, and you configure a SAML external applications SSO policy to determine the users and resources to which the SAML SSO experience applies.

To configure a SAML External Apps SSO resource policy:

1. Select **Users > Resource Policies > Web**.
2. Use the tabs to display the SSO > SAML External Apps SSO page.
3. If your administrator view is not configured to show SAML policies, click the Customize button in the upper-right corner of the page and select the SSO and SAML check boxes.
4. Click **New Policy** to display the configuration page.
5. Complete the settings described in the following table.
6. Click **Save Changes**.

The following table lists the SAML SSO External Applications Policy Configuration Guidelines:

Settings	Guidelines
Name	Type a name for the policy.
Description	Type a description that would be meaningful to other administrators.
Resources	Specify the fully qualified domain name for the resources for which this policy applies. These are the resources protected at the SAML service provider.

Settings	Guidelines
Roles	<p>Select one of the following options:</p> <p>Policy applies to ALL roles. To apply this policy to all users.</p> <p>Policy applies to SELECTED roles. To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.</p> <p>Policy applies to all roles OTHER THAN those selected below. To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.</p>
Action	<p>Select one of the following actions:</p> <p>Use the SAML SSO defined below. Typically, this is the setting you use for a SAML SSO resource policy. The system SAML identity provider makes the SSO request when a user tries to access to a SAML resource specified in the Resources list.</p> <p>Do NOT use SAML. The system does not perform an SSO request. Use this if there is a problem with the SAML service provider and you want to allow access.</p> <p>Use Detailed Rules. Use this option to configure advanced rules.</p>
SAML SSO Details	
Service Provider Entity ID	Select the service provider entity ID. The service provider entity IDs listed here are configured on the Authentication > Signing In > Sign-in SAML > Identity Provider > Peer Service Provider pages.

Configuring a SAML 2.0 ACL Web Policy

To configure the system as a policy enforcement point, you must create a SAML ACL web policy.

To configure a SAML ACL web policy:

1. In the admin console, select **Users > Resource Policies > Web**.
2. Use the tabs to display the **Access > SAML ACL page**.

If your administrator view is not configured to show SAML policies, click the Customize button in the upper-right corner of the page and select the SAML ACL check box.

3. On the SAML Access Control Policies page, click **New Policy**.
4. Complete the settings described in the following table.

5. Click **Save Changes**.
6. On the SAML Access Control Policies page, order the policies according to how you want the system to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

The following table lists the SAML ACL Web Policy Settings:

Setting	Description
Name	Type a name for the policy.
Description	Type a description that would be meaningful to other administrators.
Resources	Specify the fully qualified domain name for the resources for which this policy applies. These are the resources protected at the SAML service provider.
Roles	Select one of the following options: Policy applies to ALL roles. To apply this policy to all users. Policy applies to SELECTED roles. To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list. Policy applies to all roles OTHER THAN those selected below. To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
Action	Select one of the following options: Use the SAML Access Control checks defined below. The system performs an access control check to the specified URL using the data specified in the SAML Access Control Details section. Do not use SAML Access. The system does not perform an access control check. Use Detailed Rules. Use this option to configure advanced rules.

Setting	Description
SAML Access Control Details	<p>SAML Version. Select 2.0.</p> <p>Configuration Mode. If you select manual, complete the SAML Access Control details. If you select Metadata, select the policy decision point to use.</p> <p>If the metadata option is disabled, you have not defined or uploaded a metadata file on the System > Configuration > SAML page.</p> <p>SAML Web Service URL. Completed automatically if using metadata. If you configure manually, enter the URL of the access management system SAML server. For example, enter https://hostname/ws.</p> <p>SAML Web Service Issuer. Enter the hostname of the issuer, typically the hostname of the access management system.</p> <p>You must enter unique string that the SAML Web service uses to identify itself in authorization assertions.</p>
Authentication Type	<p>Select one of the following options:</p> <p>None-Do not authenticate the system.</p> <p>Username-Authenticate using a username and password. Enter the username and password that the system must send the Web service.</p> <p>Certificate Attribute-Authenticate using a certificate signed by a trusted certificate authority. If you have more than one certificate installed on the system, use the drop-down list to select which certificate to send to the Web service.</p>
User Identity	<p>Subject Name Type-Specify which method the system and SAML Web service should use to identify the user. Select one the following options:</p> <p>DN-Send the username in the format of a DN (distinguished name) attribute.</p> <p>Email Address-Send the username in the format of an e-mail address.</p> <p>Windows-Send the username in the format of a Windows domain qualified username.</p> <p>Other-Send the username in another format agreed upon by the system and the SAML Web service.</p> <p>Subject Name-Use variables to specify the username to the SAML Web service. Or, enter static text.</p> <p>You must send a username or attribute that the SAML Web service will recognize.</p>
	<p>Device Issuer-Enter a name that uniquely identifies the SAML authority, such as the device hostname.</p>

Setting	Description
Maximum Cache Time	You can eliminate the overhead of generating an authorization decision each time the user requests the same URL by indicating that the system must cache the access management system's authorization responses. Enter the amount of time the system should cache the responses (in seconds).
Ignore Query Data	By default, when a user requests a resource, the system sends the entire URL for that resource (including the query parameter) to the SAML Web service and caches the URL. You can specify that the system should remove the query string from the URL before requesting authorization or caching the authorization response.

Example: Implementing SAML 2.0 Web Browser SSO for Google Apps


This example shows how to implement SAML 2.0 Web browser SSO for Google Apps.

Topology

When deployed to support access to external resources (for example, public cloud resources), the system does not have to be a gateway to user access. The user can access the external resource directly, and the traffic does not flow through the device. You configure the system as a SAML identity provider to correspond with the external SAML service provider, and you configure a SAML SSO external applications policy to determine the users and resources to which the SAML SSO experience applies.

When you configure the SAML identity provider, some settings are necessary to support either identity-provider-initiated or service-provider-initiated SSO. The documentation for the configuration steps makes note of these settings. Regardless, you configure the SAML identity provider to support both identity-provider-initiated and service-provider-initiated SSO.

The following figure illustrates the flow of network communication in a service-provider-initiated SSO scenario:



Ivanti Connect Secure as a SAML Identity Provider (Peer Mode) in a Service-Provider-Initiated SSO Scenario:

1 - The user clicks a link to access a resource.

2a - The service provider sends an HTTP redirect status code (HTTP 302) to the user. The SAML request and all other SAML details are sent as URL parameters in the URL Location header.

2b - The user sends an HTTP GET request to the identity provider. The SAML request and all other SAML details are sent as URL parameters.

If the user already has a session with the identity provider, steps 2.1 and 2.2 are skipped.

2.1 - If the user does not have a session, the identity provider sends an authentication challenge to the user.

2.2 - The user enters sign-in credentials.


3a - The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.

3b - The user sends the form to the service provider.

4 - The external resource is delivered to the user's browser.

The following figure illustrates the flow of network communication in an identity-provider-initiated SSO scenario.

Ivanti Connect Secure as a SAML Identity Provider (Peer Mode) in an Identity-Provider-Initiated SSO Scenario:



1 - The user authenticates to the identity provider.

2 - The identity provider returns a portal page with links to external resources.

3 - The user clicks a link for an external resource.

4a - The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.

4b - The user sends the form to the service provider.

5 - The external resource is delivered to the user's browser.

Configuring the Google Apps SAML Service Provider

To configure the Google Apps SAML service provider:

1. Log into the Google Apps control panel. The URL is similar to the following:
<https://www.google.com/a/cpanel/acmegizmo.com>.
2. Click **Advanced Tools** in the menu bar.
3. Click the Set up single sign-on (SSO) link to display its configuration page, as shown in the following figure.
4. Configure the SAML service provider settings as described in the following table.
5. Click **Save Changes**.

The following figure depicts the Google Apps Advanced Tools: SSO:

Dashboard | Organization & users | Groups | Domain settings | Reports | Advanced tools | Setup | Support | Settings | Help

Your settings have been saved.

[← Back to Advanced tools](#)

Set up single sign-on (SSO)

To set up SSO, please provide the information below. [SSO Reference](#)

Enable Single Sign-on dana-na/auth/saml-ss0.cgi

Sign-in page URL *
 URL for signing in to your system and Google Apps

Sign-out page URL *
 URL to redirect users to when they sign out

Change password URL *
 URL to let users change their password in your system; when defined here, this URL is shown even when Single Sign-on is not enabled

Verification certificate *
 A certificate file has been uploaded-[Replace certificate](#)

The certificate file must contain the public key for Google to verify sign-in requests. [Learn more](#)

Use a domain specific issuer

This must be checked if your domain uses an IDP Aggregator to handle SAML requests.
 If enabled, the issuer value sent in the SAML request will be **google.com/a/saml-ss0-example.com** instead of simply **google.com** [Learn more](#)

Network masks

Network masks determine which addresses will be affected by single sign-on. If no masks are specified, SSO functionality will be applied to the entire network.
 Use a semicolon to separate the masks. Example: (64.233.187.99/8; 72.14.0.0/16)
 For ranges, use a dash. Example: (64.233.167-204.99/32)
 All network masks must end with a CIDR. [Learn more](#)

The following table lists the Google Apps SSO Configuration:

Settings	Guidelines
Enable Single Sign-on	Select this option.
Sign-in page URL	Type the URL of the system SAML SSO service. The URL formed with the primary host FQDN for SAML has the following form: https://SAMLHostName/dana-na/auth/saml-ss0.cgi For example: https://i5.lab.pulsesecure.net/dana-na/auth/saml-ss0.cgi The URL formed with the alternate host FQDN for SAML (to support Ivanti/NC session detection) has the following form: https://i5pulse.lab.pulsesecure.net/dana-na/auth/saml-ss0.cgi

Settings	Guidelines
Sign-out page URL	We recommend using the URL for the sign-in page for the realm associated with the system SAML identity provider. Users who already have a session will be directed to the sign-in page and can decide whether to log out from the system or not. The default sign-in URL has the form: https://FQDN For example: https://i5.lab.pulsesecure.net/
Change password URL	We recommend using the URL for the sign-in page for the realm associated with the system SAML identity provider. The system provides password management capabilities for some back-end auth servers (such as AD, LDAP, or Local Auth). When implemented, the password management capabilities are accessed from the sign-in page. The default sign-in URL has the form: https://FQDN For example: https://i5.lab.pulsesecure.net/
Verification certificate	Click Browse and select the device certificate. Then click Upload and ensure that the certificate is saved.
Use a domain specific issuer	Select this option.
Network masks	Do not select.


Configuring the Ivanti Connect Secure SAML Identity Provider

You configure the system SAML identity provider settings to match the Google Apps SAML service provider settings.

To configure the SAML identity provider settings:

1. Select **System > Configuration > SAML > Settings to complete the global SAML** settings. These settings apply to all of your SAML deployments. The following figure shows an example of SAML global settings.

The following figure depicts the SAML Global Settings:



SAML >
Settings

▼ **Metadata Server Configuration**

Timeout value for metadata fetch request: seconds

Validity of uploaded/downloaded metadata file: days

Host FQDN for SAML:

Alternate Host FQDN for SAML:

2. Select **Authentication > Signing In > Sign-In SAML > Identity Provider** to configure SAML identity provider settings. These settings apply to all of your deployments where the device is a SAML identity provider. [The following figure depicts the SAML Identity Provider Settings:](#) shows an example of SAML identity provider settings.

The following figure depicts the SAML Identity Provider Settings:

Signing In

Sign-in Policies | Sign-in Pages | Sign-in Notifications | **Sign-in SAML**

Metadata Provider: **Identity Provider**

Basic Identity Provider (IdP) Configuration (Published in Metadata)

Protocol Binding to use for SAML Response

Post
 Artifact

Signing Certificate: **No Signing** Certificate to use for signing SAML messages sent by this IdP

Decryption Certificate: **No Encryption** Certificate to use for decrypting the encrypted data in SAML messages sent by the Peer Service Provider (SP). This certificate is used by the peer SP to encrypt the SAML messages

Other Configurations

Reuse Existing NC (Pulse) Session If enabled, the user's existing NC (Pulse) session if any will be used in the SP-initiated SSO scenario, instead of authenticating the user again. Can be disabled in Peer SP configuration.

Accept unsigned AuthnRequest Individual SPs can choose to accept unsigned AuthnRequest.

Service-Provider-related IdP Configuration

The following settings apply to all Service Providers by default. Can be overridden in Peer SP configuration

Relay State: 'RelayState' sent to SP in IdP-initiated SSO scenario. If left blank, the (URL) identifier of the resource being accessed is sent as 'RelayState'

*Session Lifetime: None Suggested maximum duration of the session at the SP created due to SAML SSO.
 Role Based
 Customize

*SignIn Policy: **No SP-initiated SSO** The SignIn Policy used by this IdP to authenticate the user in SP-initiated SSO scenario.

*Force Authentication Behavior: Reject AuthnRequest SA behavior if SP sends an authentication request with ForceAuthn set to true for a user with valid browser session. Prevails over Pulse session use setting.
 Re-Authenticate User

User Identity

*Subject Name Format: **DN** Format of 'NameIdentifier' field in generated Assertion.

*Subject Name: **uid=<USERNAME>** Template for generating user's identity as sent in 'NameIdentifier' field.

Web Service Authentication

*Authentication Type: None Method used to authenticate the SP's assertion consumer service to the IdP. For Certificate based authentication the Client CA of the SP should be present in **Trusted Client CAs**
 Username/Password
 Certificate

Artifact configuration

*Source ID: **cltP464AGNn4Om4GCAeknn1e52Q=** 20-byte device identifier that identifies the artifact resolution service on the IdP (https://<Devicehostname>/dana-ws/saml20.ws).
 Enable Artifact Response Signing and Encryption If checked, Artifact response will be signed and encrypted.

Attribute Statement Configuration

Attributes to be sent in SAML Attribute Statements can be configured as name-value pairs and/or to be fetched from a Directory server.

Name-Value based configuration, here values can be system variables available in SSO parameter fields:

Delete

<input type="checkbox"/>	Attribute Name	Friendly Name	Attribute Value	Value Type	
<input type="checkbox"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	Single-Valued	<input type="button" value="Add"/>

The value can contain text characters as well as variable substitution. Variables should be enclosed in angle brackets <variable>.
 Examples:
 <USER> The user's login name
 <REALM> The user's sign-in realm
 <ROLE SEP=","> For a comma-separated list of roles.
 Mark an element as Multi-Valued if the attribute should be sent as a multi-valued attribute, in this case multiple values in an attribute value will be individual values

Directory server based configuration:
 Directory Server: **None** Select a directory server. Visit the **Servers** page to create or manage Directory servers.

- On the SAML Sign-In Identity Provider page, click Add SP and complete the settings for communication with Google Apps. Google Apps does not publish metadata, so the configuration is manually configured. The Google SAML service provider uses the HTTP POST binding and takes usernames in e-mail address format. The following figure shows an example of SP settings for Google Apps.

The following figure depicts the Peer SP Settings: (**Authentication > Signing In > New Peer Service Provider**)

Signing In > New Peer Service Provider

New Peer Service Provider

*Configuration Mode: Manual Metadata If metadata is selected, uses metadata files uploaded/added at [Peer SAML Metadata Providers](#).

▼ **Service Provider Configuration**

*Entity Id: Unique SAML Identifier of the SP.

*Assertion Consumer Service URL: URL of the service on SP that receives the assertion/artifact generated by the IdP.

Protocol Binding supported by the Assertion Consumer Service at the SP.

Post
 Artifact

*Default Binding: Post Artifact

Signature Verification Certificate: This certificate is used by IdP to verify the signature in the incoming SAML Messages. If no certificate is specified here then the certificate in the incoming message is used to verify the signature.

Issued To:
 Issued By:
 Valid:
 Details: ▶ Other Certificate Details

Upload Certificate: No file chosen

Encryption Certificate: The certificate to use if the the assertions from this IdP need to be encrypted.

Issued To:
 Issued By:
 Valid:
 Details: ▶ Other Certificate Details

Upload Certificate: No file chosen

▼ **Certificate Status Checking Configuration**

Enable signature verification certificate status checking Check this to enable revocation checks for the signing certificate. (Uses configuration in [Trusted Client CAs](#).)

Enable encryption certificate status checking Check this to enable revocation checks for the Encryption certificate. (Uses configuration in [Trusted Client CAs](#).)

▼ **Customize IdP Behavior**

Override Default Configuration

Reuse Existing NC (Pulse) Session If enabled, the user's existing NC (Pulse) session if any will be used in the SP-initiated SSO scenario, instead of authenticating the user again.

Accept unsigned AuthnRequest

Relay State: 'RelayState' sent to SP in IdP-initiated SSO scenario. If left blank, the (URL) identifier of the resource being accessed is sent as 'RelayState'.

*Session Lifetime: None Role Based Customize Suggested maximum duration of the session at the SP created due to SAML SSO.

*SignIn Policy: The SignIn Policy used by this IdP to authenticate the user in SP-initiated SSO scenario.

*Force Authentication Behavior: Reject AuthnRequest Re-Authenticate User SA behavior if SP sends an authentication request with ForceAuthn set to true for a user with valid browser session. Prevails over Pulse session setting.

User Identity

*Subject Name Format: Format of 'NameIdentifier' field in generated Assertion.

*Subject Name: Template for generating user's identity as sent in 'NameIdentifier' field.

Web Service Authentication

*Authentication Type: None Username/Password Certificate Method used to authenticate the SP's assertion consumer service to the IdP. For Certificate based authentication the Client CA of the SP should be in [Trusted Client CAs](#)

Artifact configuration

*Source ID: 20-byte device identifier that identifies the artifact resolution service on the IdP (<https://<Devicehostname>/dana-ws/saml20.ws>).

Enable Artifact Response Signing and Encryption If checked, Artifact response will be signed and encrypted.

Attribute Statement Configuration

Send Attribute Statements If checked, Attribute statements will be sent for the SP.

Use IdP Defined Attributes
 Customize IdP Defined Attributes

- Select Users > Resource Policies > Web > SAML External Apps SSO and complete settings for the external applications policy that controls the users and the resources that can use the SSO implementation. The following figure shows an example of an SAML external applications SSO policy for Google Apps.

The following figure depicts the SAML External Apps SSO Policy Settings:

Resource Policies > SAML External Apps SSO Policies > New Policy

New Policy Customize

* Name: Required: Label to reference this policy.

Description:

▼ Resources

Specify the resources for which this policy applies, one per line. In order for your resource comparisons to work effectively, you must enter a fully qualified domain name in your resource.
NOTE: This does not support IPv6.

* Resources:
 Examples:
*.domain.com/public
http://www.domain.com:8080/*
10.10.10.10/255.255.255.0:80
10.10.10.10/24:8000-9000

▼ Roles

Policy applies to ALL roles
 Policy applies to SELECTED roles
 Policy applies to all roles OTHER THAN those selected below

Available roles:

Add ->
 Remove

Selected roles:

▼ Actions

Use the SAML SP defined below
 Do not use SAML SP
 Use Detailed Rules(available after you click 'Save Changes')

SAML SSO Details

Service Provider Entity ID:

Save Changes Save as Copy

Verifying the Google Apps SAML SSO Deployment

Access a Google Docs or Google Apps resource as a non-admin user to verify the solution works as expected.



Use a browser plugin such as HTTP Watch if you want to trace the SAML communication between the SAML service provider and SAML identity provider.

To verify service-provider-initiated SSO:

1. Make sure you are not logged into the device or Google.
2. Open a Web browser and open a location on Google Docs or Google Apps. Google Apps redirects you to the sign-in page to authenticate.
3. Log in.

The access management framework processes the authentication request, performs host checking rules and role mapping rules. If authentication is successful, the system redirects you to the Google Docs or Google Apps location you had requested.

To verify Ivanti/NC session detection for service-provider-initiated SSO:

1. Make sure you are not logged into the device or Google.
2. Use Pulse or VPN tunneling client to create an SSL VPN connection.
3. Open a Web browser and open a location on Google Docs or Google Apps.

You should not have to authenticate to access the Google Docs or Google Apps location.

To verify identity-provider-initiated SSO:

1. Use the system admin console to create a bookmark to a location on Google Docs or Google Apps.
2. As a user, log in to the device.
3. Click the bookmark link to the Google Docs or Google Apps location.

You should not have to authenticate to access the Google Docs or Google Apps location.

Using SAML AuthnContext Class Variables in Role Mapping and Web ACL Rules

This topic describes how to use Security Assertion Markup Language (SAML) AuthnContext class variables in access management framework rules. For information about SAML AuthnContext class variables, refer to the SAML 2.0 OASIS Authn Context specification.

Configuring SAML AuthnContext Class Variables in the Authentication Server Configuration

In deployments where the system is a SAML service provider (SAML SP), you can configure the SAML SP to request authentication context classes from the SAML identity provider (SAML IdP). The SAML SP includes these in the RequestedAuthnContext element. In response, the SAML IdP sends the context data along with the authentication results.

The system stores the authnContext data in the session cache. You can use the system variable named `samlAuthnContextClass` to create rules based on AuthnContext in role mapping and resource policies.

To specify the SAML AuthnContext class variables in the SAML SP configuration:

1. Select **Authentication > Auth. Servers**.
2. Create a **new SAML server configuration** or edit one you have already created.

The following figure shows the SAML server configuration page. Red boxes highlight the configuration elements for AuthnContext classes.

3. Select the **AuthnContext** classes that you want to request from the SAML IdP, and select a comparison method.

This feature supports all authentication context classes described in the SAML 2.0 OASIS Authn Context specification.

The comparison method values are defined in the SAML 2.0 OASIS core specification. You should specify the same values that have been configured on the SAML IdP. If none is specified in the SAML IdP configuration, the implicit default is exact.

4. Save the configuration.

The following figure depicts the Authentication Server Configuration Page:

Auth Servers > New SAML Server

New SAML Server

Server Name:

▼ Settings

*SAML Version: 1.1 2.0

*Connect Secure Entity Id: Unique SAML identifier of the SAML Auth Server. Uses host name configured at SAML Settings.

*Configuration Mode: Manual Metadata Uses metadata files configured at SAML Metadata for metadata file based configuration.

*Identity Provider Entity Id: Unique SAML identifier of the Identity Provider.

Identity Provider Single Sign On Service URL: User is redirected to this URL in destination first scenario.

User Name Template:
 Example: <assertionNameDN.uid>, uid from X509SubjectName.
 The entire assertion name identifier if not specified; Or
 <userAttr.attr>, attr from AttributeStatement attributes.

Allowed Clock Skew (minutes): 0 - 9999 minutes

Support Single Logout If checked, Connect Secure supports sending and receiving single logout requests.

▼ SSO Method

Artifact Post

Response Signing Certificate:
 Issued To:
 Issued By:
 Valid:
 Details: ▶ Other Certificate Details

Upload Certificate: No file chosen

Enable Signing Certificate status checking
 (Uses configuration in Trusted Client CAs. This applies to the certificate configured above as well as the one comes along with the SAML response.)

Select Device Certificate for Signing: Certificate used for signing the Requests initiated by Connect Secure for the SAML Auth Server. Applicable if Request signing is not required.

Select Device Certificate for Encryption: Certificate used by the IdP for wrapping encryption keys for the SAML Auth Server. Select "Not" if encryption is not required.

Select Requested Authn Context Classes to be sent in the AuthRequest:

Available:	Selected:
InternetProtocolPassword	PasswordProtectedTransport
Kerberos	
MobileOneFactorUnregistered	
MobileTwoFactorUnregistered	
X509	

Comparison Method for Authentication Classes:

▼ Service Provider Metadata Settings

Metadata Validity: days 1 - 9999. Specifies the time in days after which metadata for the SAML Auth Server should be refreshed by the Identity Provider. This is used to populate the cache in the generated metadata.

Do Not Publish Connect Secure Metadata Prevents the Metadata for the SAML Auth Server to be published at the location specified by the Connect Secure Entity Id.

▼ User Record Synchronization

Enable User Record Synchronization

Logical Auth Server Name:

Configuring a Role Mapping Rule Based on a SAML AuthnContext Class Variable

You can use role mapping rule custom expressions to include AuthnContext class data as a factor in role determination.

To configure role mapping rules:

1. Select **Users > User Realms**.
2. Create a new realm or edit a realm you have already created.
3. Click **New Rule** to display the configuration page.
4. Select Custom Expression and click Update to redisplay the configuration page with the controls related to custom expressions.

The following figure shows the configuration page.

5. Click Expressions to display the server catalog dialog box.

The next figure shows the dialog box.

6. Select `samlAuthnContextClass`, select an operator, and click Insert Expression.
7. Edit the expression template to match the `AuthnContextClassRef` data expected from the SAML IdP.
8. Save your changes to the variable expression and return to the rule configuration page.
9. Select the expression, roles for the rule, and the stop option (if desired).
10. Save your changes to the rule configuration and return to the realm configuration page.
11. Reorder the rules if necessary.
12. Save the realm configuration.

The following figure depicts the Role Mapping Rule Configuration Page:

User Realms > SAMLRealm > Role Mapping > Role Mapping Rule

Role Mapping Rule

Rule based on: Custom Expressions

* Name:

▼ Rule:if user has any of these custom expressions...

Available Expressions: (none)

Selected Expressions: ProtectedTransport

▼ then assign these roles

Available Roles: Android_CloudSecure_Role
 CloudSecure_Removed_Role
 IOS_CloudSecure_Role
 Mac_CloudSecure_Role
 Users
 Windows_CloudSecure_Role

Selected Roles: FullAccess

Stop processing rules when this rule matches

To manage roles, see the [Roles](#) configuration page.

The following figure depicts the Server Catalog Expressions and Variables:

Server Catalog for Consumer

Expressions Variables

View: ProtectedTransport

Name:

Expression: ProtectedTransport"/>

Expressions Dictionary

- ▶
- ▶
- ▶
- ▶
- ▶
- ▶
- ▶
- ▶
- ▶
- ▶

Configuring a Web ACL Policy Rule Based on a SAML AuthnContext Class Variable

You can use the resource policy detailed rules configuration to include AuthnContext class data as a factor in resource access determinations. This example shows how to use a SAML AuthnContext class variable in Web ACL detailed rules. In the same manner, you can use the AuthnContext class variable in detailed rules for other resource policies.

To configure a resource policy:

1. Select Resource **Policies > Web > Web ACL**.
2. Create a new policy or edit a policy you have already created.
3. Click the **Detailed Rules** tab for the policy.
4. Click **New Rule** to display the detailed rules configuration page.

The following figure shows the detailed rule configuration page.

5. Under Conditions, select **samlAuthnContextClass**, select an operator, and click Insert Expression.
6. Edit the condition expression template to match the **AuthnContextClassRef** data expected from the SAML IdP.
7. Select a rule action and resources to which the rule applies, and save your changes to return to the policy configuration page.
8. Reorder the rules if necessary.
9. Save the configuration.

The following figure depicts the Detailed Rule Configuration Page:

Web Access Policies > Initial Policy for Local Resources

Detailed Rule

Actions

Allow access
 Deny access

Resources

Specify the resources for which this rule applies, one per line.

*Resources:

Examples:
 http://*.domain.com/public/*
 https://www.domain.com:443/*
 10.10.10.10/255.255.255.0:80,443/public/*
 10.10.10.10/24:8000-9000/*

Conditions

Specify the conditions, if any, under which this rule applies.

Conditions:

[Click here to save the above condition in the catalog.](#)

Conditions Dictionary

- samlAuthnContextClass
- samlMultiValAttr.<auth-attr>
- sourceIp
- time
- time.day
- time.dayOfWeek
- time.dayOfYear
- time.month
- time.year
- user

=

[< Insert Expression](#)

[Save Changes](#) [Save as Copy](#)

Using Policy Tracing Logs to Verify the SAML AuthnContext Class Variable Is Used in Rules

You can use policy tracing logs to verify your configuration.

To create a policy trace log:

1. Select **Troubleshooting > User Sessions > Policy Tracing** to display the configuration page.
2. Specify the **username**, **realm**, and **source IP** address if you know it. If you provide the source IP address, the policy trace log can include events that occur before the **user ID** is entered into the system.
3. Select the events to trace.

4. Click **Start Recording**.
5. Initiate the action you want to trace, such as a user sign in.
6. Click **View Log** to display the policy trace results log.
7. Click **Stop Recording** when you have enough information.

The following figure shows policy trace results. The highlighted entries show the data populated to the `samlAuthnContextClass` system variable, as well as the role mapping rule that was matched.

Current Policy Trace Log		
Date:	Earliest Date to Latest Date	
User Name:	jumbo	
Realm Name:	SAMLRealm	
Export Format:	Standard	
Show	1000	items
<input type="button" value="Update"/> <input type="button" value="Save Log As..."/> <input type="button" value="Clear Log"/>		
Severity	ID	Message
Info	PTR23344	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Authentication successful to auth server "Consumer"
Info	PTR10209	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Realm SAMLRealm running 2 mapping rules for user jumbo
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable sourceIp = 10.206.152.145
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable SourceIPStr = "10.206.152.145"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable user = "jumbo"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable password = "*****"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable userName = "jumbo"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable protocol =
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable realm = "SAMLRealm"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable loginTime = Thu Oct 31 15:17:18 2013
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable userAgent = "Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 6.1; WOW64; Trident/4.0; SLCC2; .NET CLR 2.0.50727; .NET CLR 3.5.30729; .NET CLR 3.0.30729; Media Center PC 6.0; .NET4.0C)"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable language = "en"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable loginURL = "*/saml/"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable loginHost = "samlconsumer.qalab.com"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable loginHostAddr = "10.204.55.40"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable assertRef = "e97d9fb8920635a2f4c13c989b348bfa"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable networkIF = "internal"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable SessionIndex = ""
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable nameIdFormat = "urn:oasis:names:tc:SAML:1.1:nameid-format:X509SubjectName"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable samlName = "uid=jumbo"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable samlNameQ = ""
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable samlAuthnContextClass = "urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport"
Info	PTR10305	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Variable user@Consumer = "jumbo"
Info	PTR10212	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Mapped to roles FullAccess by rule 'samlAuthnContextClass = "urn:oasis:names:tc:SAML:2.0:ac:classes:PasswordProtectedTransport"'
Info	PTR10213	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Role mapping stopped by Stop rule
Info	PTR10205	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Realm SAMLRealm mapped user jumbo to roles FullAccess
Info	PTR23353	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[] - Role restrictions successfully passed for roles: FullAccess
Info	PTR23362	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[FullAccess] - Sign-in successful, creating session
Info	PTR23363	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[FullAccess] - Session created, redirecting user to start page. Sign-in done.
Info	PTR24559	2013/10/31 15:17:18 - [10.206.152.145] - Default Network::jumbo(SAMLRealm)[FullAccess] - Automatically redirected from page "login" to the next start page "/dana/home/index.cgi" before starting the session.

Investigating a "No valid assertion found in SAML response" Error

Problem Description: SAML is a time sensitive protocol. The time-based validity of a SAML assertion is determined by the SAML identity provider. If the SAML identity provider and SAML service provider clocks are askew, the assertion can be determined invalid, and authentication fails.

Environment: In the scenario described here, the system is deployed as a SAML service provider in a SAML 2.0 deployment.

Symptoms: In this scenario, the following error is returned to the user after the user has submitted credentials to the SAML identity provider:

```
SAML Transferred failed. Please contact your system administrator.
```

```
Detail: Failure: No valid assertion found in SAML response."
```

Cause: To investigate the error:

1. Select Maintenance > Troubleshooting > Monitoring > Debug Logs to display the Debug Log configuration page, shown in the following figure.

The following figure depicts the Debug Log Page:

The screenshot shows the 'Debug Log' configuration page. The breadcrumb trail is 'Troubleshooting > Monitoring > Debug Log'. The page title is 'Debug Log'. There are tabs for 'User Sessions', 'Monitoring' (selected), 'Tools', 'System Snapshot', and 'Remote Debugging'. Below these are sub-tabs for 'Debug Log' (selected), 'Node Monitor', 'Cluster', and 'Diagnostic Logs'. There are four buttons: 'Save Changes', 'Reset', 'Save Debug Log', and 'Clear Log...'. Under 'Debug Log Settings', the following settings are visible: 'Current Log Size' is 3828610 bytes; 'Debug Logging On' is checked; 'Max Debug Log Size' is 2 MB; 'Debug Log Detail Level' is 10; 'Include logs' is checked; 'Process Names' is empty; and 'Event Codes' is 'saml'. Red boxes highlight the 'Debug Logging On' checkbox, the '10' in the 'Debug Log Detail Level' input, and the 'saml' in the 'Event Codes' input.

2. Turn debug logging on, set Debug Log Detail Level to 10, and Event Codes to saml.

3. Reproduce the action that results in the error-in this case, user access to the resource associated with the SAML service provider that prompts the user to submit credentials to the SAML identity provider.
4. Click **Save Debug Log**.

The console displays the Save As dialog box.

5. Save the file to a location your local host or a location that you can access when sending mail. The file is an encrypted file, so do not try to open it and analyze it yourself.
6. E-mail the debug log to Global Support Center.

Global Support Center will use the file to diagnose the issue. In the debug log, the following log lines indicate issues with the time-based validity of the assertion:

```
verifySubjectConfirmationData: assertion has expired
processConditions: assertion has expired [NotOnOrAfter condition failed]
processConditions: assertion is not yet Valid [NotBefore condition failed]
```

These log lines indicate a clock sync issue only if failure of the time-based validity check is unexpected. The same log lines might appear in the debug log to indicate an assertion has expired as expected.

Solution	We recommend you use NTP to ensure the clocks are synchronized and that you set an Allowed Clock Skew value that accommodates any expected or permissible skew. Properly synchronized clocks avoid unexpected failure.
----------	--

To configure NTP:

1. Select **System > Status** to display the System Status page.
2. Next to **System Date & Time**, click **Edit** to display the Date and Time page.
3. Specify the settings for the same NTP server used by the SAML identity provider.
4. Save your configuration.



To be NDcPP compliant, NTP Update Interval needs to be limited to 60 minutes. This is to avoid the potential drift becoming too excessive.

To set the Allowed Clock Skew value:

1. Select **Authentication > Auth. Servers**.

2. Select the **SAML authentication server** you want to configure to display its configuration page.
3. Specify a number of minutes in the **Allowed Clock Skew** to accommodate any expected or permissible skew.
4. Save your configuration.

Ivanti Connect Secure SAML 1.1 Support

The trend in SAML deployments is converging on the SAML 2.0 specification. Ivanti Connect Secure continues to support SAML 1.1. The following sections reprint previous information we have provided about SAML 1.1 deployments:

- [About SAML Version 1.1](#)
- [SAML Version 1.1 Configuration Tasks](#)

About SAML Version 1.1

The following topics provide background information about SAML version 1.1:

- [Understanding SAML 1.1](#)
- [Understanding SAML 1.1 Profiles](#)
- [Understanding SAML 1.1 Assertions](#)
- [Creating a Trust Relationship Between SAML 1.1 Systems](#)

Understanding SAML 1.1

The system enables you to pass user and session state information between the device and another trusted access management system that supports the Security Assertion Markup Language (SAML). SAML provides a mechanism for two disparate systems to create and exchange authentication and authorization information using an XML framework, minimizing the need for users to re-enter their credentials when accessing multiple applications or domains.

SAML exchanges are dependent upon a trusted relationship between two systems or domains. In the exchanges, one system acts as a SAML authority (also called an asserting party or SAML responder) that asserts information about the user. The other system acts as a relying party (also called a SAML receiver) that relies on the statement (also called an assertion) provided by the SAML authority. If it chooses to trust the SAML authority, the relying party authenticates or authorizes the user based on the information provided by the SAML authority.

The system supports two SAML use case scenarios:

- The system as the SAML authority-The user signs into a resource by way of the device first, and all other systems are SAML receivers, relying on the system for authentication and authorization of the user. Under this scenario, the system can use either an artifact profile or a POST profile.
- The system as the SAML receiver-The user signs into another system on the network first, and the system is the SAML receiver, relying on the other system for authentication and authorization of the user.

For example, in the first scenario, an authenticated user named John Smith may try to access a resource protected by an access management system. When he does, the system acts as a SAML authority and declares "This user is John Smith. He was authenticated using a password mechanism." The access management system (the relying party) receives this statement and chooses to trust the system (and therefore trust that the system has properly identified the user). The access management system may still choose to deny the user access to the requested resource (for instance, because John Smith has insufficient access privileges on the system), while trusting the information sent by the system.

In the second scenario, John Smith signs in to his company portal and is authenticated using an LDAP server sitting behind the company's firewall. On the company's secure portal, John Smith clicks a link to a resource protected by the system. The following process occurs:

1. The link redirects John Smith to an intersite transfer service on the company portal, which constructs an artifact URL. The artifact URL contains a reference to a SAML assertion stored in the company portal's cache.
2. The portal sends the URL to the system, which can decide whether or not to link to the reference.
3. If the system links to the reference, the portal sends a SOAP message containing the SAML assertion (an XML message containing the user's credentials) to the system, which can then decide whether or not to allow the user access to the requested resource.



SOAP requests generated by the system (when configured as a SAML 1.1 consumer) are not signed.

- If the system allows the user access, the system presents to the user the requested resource.
- If the system rejects the SAML assertion, or the user credentials, the system responds to the user with an error message.

When configuring the system, you can use SAML for:

- **Single sign-on (SSO) authentication**-In a SAML SSO transaction, an authenticated user is seamlessly signed into another system without resubmitting his credentials. In this type of transaction, the system can be either the SAML authority or the SAML receiver. When acting as the SAML authority, the system makes an authentication statement, which declares the user's username and how he was authenticated. If the relying party (called an assertion consumer service in SAML SSO transactions) chooses to trust the system, the user is seamlessly signed into the assertion consumer service using the username contained in the statement.

When acting as the SAML receiver, the system requests credential confirmation from the SAML authority, which is the other access management system, such as LDAP or another authentication server. The SAML authority sends an assertion by way of a SOAP message. The assertion is a set of XML statements that the system must interpret, based on criteria that the system administrator has specified in a SAML server instance definition. If the system chooses to trust the asserting party, the system allows the user to sign in seamlessly using the credentials contained in the SAML assertion.

- **Access control authorization**-In a SAML access control transaction, the system asks an access management system whether the user has access. In this type of transaction, the system is the relying party (also called a policy enforcement point in access control transactions). It consumes and enforces an authorization decision statement provided by the access management system (SAML authority), which declares what the user is allowed to access. If the SAML authority (also called a policy decision point in access control transactions) declares that the user has sufficient access privileges, the user may access the requested resource

The system does not generate authorization decision statements-it only consumes them.

In addition to providing users access to a URL based on the authorization decision statement returned by a SAML authority, the system also allows you to define users' access rights to a URL using system-only mechanisms (Users > Resource Profiles > Web Applications/Pages tab). If you define access controls through the system as well as through a SAML authority, both sources must grant access to a URL for a user to access it. For example, you may configure an access policy that denies members of the "Users" role access to www.google.com, but configure another SAML policy that bases a user's access rights on an attribute in an access management system. Even if the access management system permits users access to www.google.com, users are still denied access based on the access policy.

When asked if a user may access a resource, access management systems that support SAML may return a response of permit, deny, or indeterminate. If the system receives an indeterminate response, it denies the user access.

The session timeouts on the system and your access management system may not coordinate with one another. If a user's access management system session cookie times out before his destination signaling identifier (DSID) cookie times out, then single sign-on between the two systems is lost. The user is forced to sign in again when he times out of the access management system.

Understanding SAML 1.1 Profiles

The system accepts authentication assertions generated by a SAML authority using either an artifact profile or a POST profile. This feature allows a user to sign in to a source site or portal without going through the system first, and then to access the system with single sign-on (SSO) through the SAML consumer service.

As a result, the user who authenticates elsewhere can access resources behind the device without signing in again.

Using the Artifact Profile and POST Profile

The two supported profiles provide different methods of accomplishing the same task. The end user's goal is to sign in to all desired resources once, without experiencing multiple sign-in pages for different resources or applications. Although the end user wants transparency, you, the administrator, want to ensure complete security across the resources on your system, regardless of the servers or sites represented.

The artifact profile requires that you construct an automated request-response HTTP message that the browser can retrieve based on an HTTP GET request.

The POST profile requires that you construct an HTML form that can contain the SAML assertion, and which can be submitted by an end user action or a script action, using an HTTP POST method.

Using the Artifact Profile Scenario

The SAML server generally supports the following artifact profile scenario:

1. The user accesses a source site through a browser. The source site might be a corporate portal using a non-Ivanti Connect Secure authentication access management system.
2. The source site challenges the user for username and password.
3. The user provides username and password, which the source site authenticates through a call to an LDAP directory or other authentication server.
4. The user then clicks a link on the source site, which points to a resource on a server that is protected behind the device.
5. The link redirects the user to the intersite transfer service URL on the source site. The source site pulls an authentication assertion message from its cache and encloses it in a SOAP message. The source site constructs a SAML artifact (a Base64 string) that it returns to the browser in a URI along with the destination and assertion address.
6. The destination site queries the authenticated assertion from the source site, based on the artifact it receives from the source site.
7. The system accepts the assertion as a valid authentication if the elapsed time falls within the allowable clock skew time. If the user also meets the other policy restrictions, the system grants the user access to the requested resource.

The main tasks you are required to fulfill to support the system as the relying party with the artifact profile include:

1. Implement the assertion consumer service, which:
 1. Receives the redirect URL containing the artifact.
 2. Generates and sends the SAML request.
 3. Receives and processes the SAML response.
2. Integrate the assertion consumer service with the existing system process, which:

1. Maps the SAML assertion to a local user.
2. Creates a user session.
3. Performs local authorization.
4. Serves the resource or denies access.

Using the POST Profile Scenario

The SAML server generally supports the POST profile scenario, as follows:

1. The end user accesses the source web site, hereafter known as the source site.
2. The source site verifies whether or not the user has a current session.
3. If not, the source site prompts the user to enter user credentials.
4. The user supplies credentials, for example, username and password.
5. If the authentication is successful, the source site authentication server creates a session for the user and displays the appropriate welcome page of the portal application.
6. The user then selects a menu option or link that points to a resource or application on a destination web site.
7. The portal application directs the request to the local intersite transfer service, which can be hosted on the source site. The request contains the URL of the resource on the destination site, in other words, the TARGET URL.
8. The intersite transfer service sends an HTML form back to the browser. The HTML FORM contains a SAML response, within which is a SAML assertion. The response must be digitally signed. Typically, the HTML FORM will contain an input or submit action that will result in an HTTP POST. This can be a user-clickable Submit button or a script that initiates the HTTP POST programmatically.
9. The browser, either due to a user action or by way of an auto-submit action, sends an HTTP POST containing the SAML response to the destination web site's assertion consumer service.
10. The replying party's assertion consumer (in this case, on the destination web site) validates the digital signature on the SAML response.
11. If valid, the assertion consumer sends a redirect to the browser, causing the browser to access the TARGET resource.

12. The system, on the destination site, verifies that the user is authorized to access the destination site and the TARGET resource.
13. If the user is authorized to access the destination site and the TARGET resource, the system returns the TARGET resource to the browser.

The main tasks you are required to fulfill to support the system as the relying party with the POST profile include:

1. Implement the assertion consumer service, which receives and processes the POST form
2. Integrate the assertion consumer service with the existing process, which:
 1. Maps the SAML assertion to a local user.
 2. Creates a user session.
 3. Performs local authorization.
 4. Serves the resource or denies access.

Understanding SAML 1.1 Assertions

Each party in the request-response communication must adhere to certain requirements. The requirements provide a predictable infrastructure so that the assertions and artifacts can be processed correctly.

- The artifact is a Base64-encoded string of 40 bytes. An artifact acts as a token that references an assertion on the source site, so the artifact holder-the Ivanti Connect Secure device-can authenticate a user who has signed in to the source site and who now wants to access a resource protected by the system. The source site sends the artifact to the device in a redirect, after the user attempts to access a resource protected by the system. The artifact contains:
 - TypeCode - A 2-byte hex code of 0x0001 that identifies the artifact type.

- SourceID - A Base64-encoded string of 20 bytes that determines the source site identity and location. You can use OpenSSL or similar Base64 encoding tool to generate the encoded string. The system maintains a table of SourceID values and the URL for the corresponding SAML responder. The system and the source site communicate this information in a back channel. On receiving the SAML artifact, the system decodes it and ensures that it is 20 bytes. It determines whether or not the SourceID belongs to a known source site, and, if it does, obtains the site location before sending a SAML request. The source site generates the SourceID by computing the SHA-1 hash of the source site's own URL.
- AssertionHandle - A 20-byte random value that identifies an assertion stored or generated by the source site. At least 8 bytes of this value should be obtained from a cryptographically secure RNG or PRNG.
- The intersite transfer service is the identity provider URL on the source site (not the Ivanti Connect Secure device). Your specification of this URL in the admin console enables the system to construct an authentication request to the source site, which holds the user's credentials in cache. The request is similar to the following example:

```
GET http://<intersite transfer hostname and path>?TARGET=<Target>...<HTTP-
Version><other HTTP 1.0 or 1.1 components>
```

In the preceding sample, <intersite transfer hostname and path> consists of the hostname, port number, and path components of the intersite transfer URL at the source and where Target= <Target> specifies the requested target resource at the destination (Ivanti Connect Secure protected) site. This request might look like:

```
GET http://10.56.1.123:8002/xferSvc?TARGET=http://www.dest.com/sales.htm
```

- The intersite transfer service redirects the user's browser to the assertion consumer service at the destination site-in this case, the Connect Secure device. The HTTP response from the source site intersite transfer service must be in the following format:

```
<HTTP-Version> 302 <Reason Phrase>
<other headers>
Location: http://<assertion consumer hostname and path>?<SAML
searchpart><other HTTP 1.0 or 1.1 components>
```

In the preceding sample, <assertion consumer hostname and path> provides the hostname, port number, and path components of an assertion consumer URL at the destination site and where <SAML searchpart> = ...TARGET=<Target> ...SAMLart=<SAML artifact>... consists of one target description, which must be included in the <SAML searchpart> component. At least one SAML artifact must be included in the SAML <SAML searchpart> component. The asserting party can include multiple SAML artifacts.



You can use status code 302 to indicate that the requested resource resides temporarily under a different URI.

If <SAML searchpart> contains more than one artifact, all of the artifacts must share the same SourceID.

The redirect might look like:

```
HTTP/1.1 302 Found
Location:
http://www.ive.com:5802/artifact?TARGET=/www.ive.com/&SAMLart=artifact
```

- The user's browser accesses the assertion consumer service, with a SAML artifact representing the user's authentication information attached to the URL.

The HTTP request must appear as follows:

```
GET http://<assertion consumer hostname and path>?<SAML searchpart> <HTTP-
Version><other HTTP 1.0 or 1.1 request components>
```

In the preceding sample, <assertion consumer hostname and path> provides the hostname, port number, and path components of an assertion consumer URL at the destination site.

<SAML searchpart> = ...TARGET=<Target>...SAMLart=<SAML artifact> ...

A single target description MUST be included in the <SAML searchpart> component. At least one SAML artifact MUST be included in the <SAML searchpart> component; multiple SAML artifacts MAY be included. If more than one artifact is carried within <SAML searchpart>, all the artifacts MUST have the same SourceID.

You should not expose the assertion consumer URL unless over SSL 3.0 or TLS 1.0. Otherwise, transmitted artifacts might be available in plain text to an attacker.

- The issuer value is typically the URL of the source site. You can specify the <ISSUER> variable, which will return the issuer value from the assertion.
- The username template is a reference to the SAML name identifier element, which allows the asserting party to provide a format for the username. The SAML specification allows for values in the following formats:
 - Unspecified - Indicates that interpretation of the content is left up to the individual implementations. In this case, you can use the variable assertionName.
 - E-mail Address - Indicates that the content is in the form of an e-mail address. In this case, you can use the variable assertionName.
 - X.509 Subject Name - Indicates that the content is in the form of an X.509 subject name. In this case, you can use the variable assertionNameDN.<RDN>.
 - Windows Domain Qualified Name - Indicates that the content is a string in the form of DomainName\Username.
 - You should define the username template to accept the type of username your SAML assertion contains.
- You can prevent eavesdropping on the SAML artifact by synchronizing the clocks on the source and destination sites. The system provides an Allowed Clock Skew attribute that dictates the maximum time difference allowed between the system and the source site. The system rejects any assertions whose timing exceeds the allowed clock skew.

Creating a Trust Relationship Between SAML 1.1 Systems

In order to ensure that SAML-enabled systems are only passing information between trusted sources, you must create a trust relationship between the applications that are sending and receiving information.

Configuring Trusted Application URLs

In a trust relationship, you must provide the SAML-enabled systems with the URLs they need to contact each other. In some transactions, only the system that initiates the transaction (the Ivanti Connect Secure device) needs to know the URL of the other system. (The system uses the URL to initiate the transaction.) In other transactions (SSO transactions using artifact profiles), you need to configure each system with the URL of the other.

The following list shows the different transaction types and the URLs you must configure for each:

- SSO transactions: Artifact profile - On Ivanti Connect Secure, you must enter the URL of the assertion consumer service. For example, use `https://hostname/acs`.

You must also enter the following URL for the system on the assertion consumer service. For example, use `https://<SecureAccessHostname>/dana-ws/saml.ws`.

- SSO transactions: POST profile - On Connect Secure, you must enter the URL of the assertion consumer service. For example, use `https://hostname/acs`.
- Access control transactions - On Connect Secure, you must enter the URL of the SAML Web service. For example, use `https://hostname/ws`.

Configuring an Issuer

Before accepting a statement from another system, a SAML-enabled entity must trust the issuer of the statement. You can control which issuers a system trusts by specifying the unique strings of the trusted issuers during the system's configuration. (When sending a statement, an issuer identifies itself by including its unique string in the statement. SAML-enabled applications generally use hostnames to identify issuers, but the SAML standard allows applications to use any string.) If you do not configure a system to recognize an issuer's unique string, the system will not accept that issuer's statements.

The following list shows the different transaction types and the issuers you must configure for each:

- SSO transactions-You must specify a unique string on the system (typically its hostname) that it can use to identify itself and then configure the access management system to recognize that string.
- Access control transactions-You must specify a unique string on the access management system (typically its hostname) that it can use to identify itself and then configure the system to recognize that string.

Configuring Certificates

Within SSL transactions, the server must present a certificate to the client, and then the client must verify (at minimum) that it trusts the certificate authority who issued the server's certificate before accepting the information. You can configure all of the system SAML transactions to use SSL (HTTPS).

Configuring SSO Transactions: Artifact Profile

Artifact profile transactions involve numerous communications back and forth between the system and the access management system. The methods you use to pass data and authenticate the two systems affect which certificates you must install and configure.

The following list shows the different artifact profile configuration options that require special certificate configurations:

- All artifact profile transactions-Regardless of your artifact profile configuration, you must install the certificate of the CA that signed the system Web server certificate on the access management system. (The system requires the access management system to use an SSL connection when requesting an authentication statement. In an SSL connection, the initiator must trust the system to which it is connecting. By installing the CA certificate on the access management system, you ensure that the access management system will trust the CA that issued the system certificate.)
- Sending artifacts over an SSL connection (HTTPS GET requests)-If you choose to send artifacts to the access management system using an SSL connection, you must install the access management system's root CA certificate on the system. (In an SSL connection, the initiator must trust the system to which it is connecting. By installing the access management system's CA certificate on the system, you ensure that the system will trust the CA that issued the access management system's certificate.) You can install the root CA from the System > Configuration > Certificates > Trusted Client CAs page in the admin console. If you do not want to send artifacts over an SSL connection, you do not need to install any additional certificates.

To enable SSL-based communications from the system to the access management system, enter a URL that begins with HTTPS in the SAML Assertion Consumer Service URL field during the system configuration. You may also need to enable SSL on the access management system.

- Transactions using certificate authentication-If you choose to authenticate the access management system using a certificate, you must:
 - Install the access management system's root CA certificate on the system. You can install the root CA from the System > Configuration > Certificates > Trusted Client CAs page in the admin console.
 - Specify which certificate values the system should use to validate the access management system. You must use values that match the values contained in the access management server's certificate.

If you do not choose to authenticate the access management system, or if you choose to use username/password authentication, you do not need to install any additional certificates.

Configuring SSO Transactions: POST Profile

In a POST profile transaction, the system sends signed authentication statements to the access management system. Generally, it sends them over an SSL connection (recommended), but in some configurations, the system may send statements through a standard HTTP connection.

The following list shows the different POST profile configuration options that require special certificate configurations:

- All POST profile transactions-Regardless of your POST profile configuration, you must specify which certificate the system should use to sign its statements. You can choose a certificate in the Users > Resource Policies > Web > SSO SAML > [Policy] > General page in the admin console. Then, you must install the device certificate on the access management system. You can download the certificate from the System > Configuration > Certificates > Device Certificates > [Certificate] > Certificate Details page.
- Sending POST data over an SSL connection (HTTPS)-If you choose to send statements to the access management system using an SSL connection, you must install the access management system's root CA certificate on the system. (In an SSL connection, the initiator must trust the system to which it is connecting. By installing the access management system's certificate on the system, you ensure that the system will trust the CA that issued the access management system's certificate.) You can install the root CA from the System > Configuration > Certificates > Trusted Client CAs page in the admin console. If you do not want to post statements over an SSL connection, you do not need to install any additional certificates.

To enable SSL-based communications from the system to the access management system, enter a URL that begins with HTTPS in the SAML assertion consumer service URL field during the system configuration. You may also need to enable SSL on the access management system.

Configuring Access Control Transactions

In an access control transaction, the system posts an authorization decision query to the access management system. To ensure that the access management system responds to the query, you must determine which certificate options are required by your configuration.

The following list shows the different access control configuration options that require special certificate configurations:

- Sending authorization data over an SSL connection-If you choose to connect to the access management system using an SSL connection, you must install the access management system's root CA on the system. (In an SSL connection, the initiator must trust the system to which it is connecting. By installing the access management system's certificate on the system, you ensure that the system will trust the CA that issued the access management system's certificate.) You can install the root CA from the System > Configuration > Certificates > Trusted Client CAs page in the admin console.

- Transactions using certificate authentication-If you choose to use certificate authentication, you must configure the access management system to trust the CA that issued the certificate. Optionally, you may also choose to accept the certificate based on the following additional options:
 - Upload the certificate public key to the access management system.
 - Validate the system using specific certificate attributes.

These options require that you specify which certificate the system should pass to the access management system. You can choose a certificate in the Users > Resource Policies > Web > SAML ACL > [Policy] > General page in the admin console.

To determine how to configure your access management system to validate the certificate, see your access management system's documentation. If your access management system does not require certificate authentication, or if it uses username/password authentication, you do not need to configure the system to pass the access management server a certificate. If you do not specify a trust method, your access management system may accept authorization requests from any system.

Configuring User Identity

In a trust relationship, the two entities must agree on a way to identify users. You may choose to share a username across systems, select an LDAP or certificate user attribute to share across systems, or hardcode a user ID. (For example, you may choose to set the Subject Name field to "guest" to easily allow access across systems.)

To ensure that the two systems are passing common information about users, you must specify which information the system should pass using options in the User Identity section of the Users > Resource Policies > Web > SAML SSO > [Policy] > General page and the Users > Resource Policies > Web > SAML ACL > [Policy] > General page. Choose a username or attribute that the access management system will recognize.

SAML Version 1.1 Configuration Tasks

The following topics describe how to configure the features that support SAML version 1.1:

- [Creating a SAML 1.1 Server Instance](#)
- [Configuring SAML 1.1 SSO Profiles](#)
- [Creating a SAML 1.1 SSO POST Profile](#)
- [Creating a SAML 1.1 ACL Resource Policy](#)

Creating a SAML 1.1 Server Instance

To create a new SAML server instance:

1. In the admin console, choose **Authentication > Auth. Servers**.

Select SAML Server from the New list, and then click New Server. Complete the settings as described in the following table.

2. Click **Save Changes**.

After you save changes for the first time, the page is redisplayed and now has two tabs. The Settings tab allows you to modify any of the settings pertaining to the SAML Server instance. The Users tab lists valid users of the server.

The following table lists the SAML Authentication Server (SAML 1

Setting	Guideline
Name	Specify a name to identify the server instance.
Settings	
SAML Version	Select 1.1.
Source Site Inter-Site Transfer Service URL	User is redirected to this URL in destination first scenario.
Issuer Value for Source Site	Typically, the URI or hostname of the issuer of the assertion.
User Name Template	Enter the mapping string from the SAML assertion to a user realm. For example, enter <assertionNameDN.CN> to derive the username from the CN value in the assertion.
Allowed Clock Skew	The maximum allowed difference in time between the system clock and the source site clock.

Setting	Guideline
	<p>SAML is a time sensitive protocol. The time-based validity of a SAML assertion is determined by the SAML identity provider. If the SAML identity provider and SAML service provider clocks are askew, the assertion can be determined invalid, and you will receive the following error:</p> <p>SAML Transferred failed. Please contact your system administrator. Detail: Failure: No valid assertion found in SAML response.</p> <p>We recommend you use NTP to ensure the clocks are synchronized and that you set an Allowed Clock Skew value that accommodates any expected or permissible skew.</p>
SSO Method	
Artifact	<p>Source ID. A Base64-encoded string of 20 bytes that the system uses to recognize an assertion from a given source site.</p> <p>Source SOAP Responder Service URL</p> <p>SOAP Client Authentication. Select HTTP Basic or SSL Client Certificate and complete the related settings.</p> <p>SOAP requests generated by the system (when configured as a SAML 1.1 consumer) are not signed.</p>
POST	<p>Response Signing Certificate. Enter the name of, or browse to locate, the PEM-formatted signing certificate, which is loaded for the SAML response signature verification.</p> <p>The certificate you select should be the same certificate used for signing the SAML response at the source site. The source site may send this certificate along with the SAML response, depending on the source site configuration. By default, the system performs signature verification of the SAML response first on the locally configured certificate. If a certificate is not configured locally in the SAML authentication server, then the system performs the signature verification on the certificate included in the SAML response from the source site.</p> <p>Enable Signing Certificate status checking. Select this option to check the validity of the signing certificate configured in the SAML authentication server POST profile. It is possible that the certificate has already expired or has been revoked.</p>
User Record Synchronization	

Setting	Guideline
Enable User Record Synchronization	Allow users to retain their bookmarks and individual preferences regardless of which device they log in to.
Logical Auth Server Name	Logical name of the authentication server.

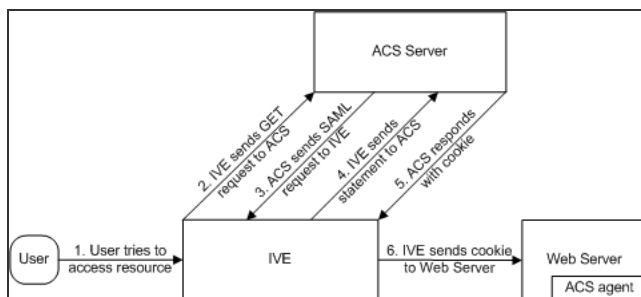
Configuring SAML 1.1 SSO Profiles

When enabling SSO transactions to a trusted access management system, you must indicate whether the access management system should "pull" user information from Connect Secure or whether Connect Secure should "push" it to the access management system. You indicate which communication method the two systems should use by selecting a profile during configuration. A profile is a method that two trusted sites use to transfer a SAML statement. When configuring the system, you may choose to use an artifact or POST profile.

When you choose to communicate using the artifact profile (also called browser/artifact profile), the trusted access management server "pulls" authentication information from the system.

The following figure shows the SAML communication process when the implementation uses the artifact profile.

The following figure depicts the Artifact Profile:



The system and an assertion consumer service (ACS) use the following process to pass information:

1. The user tries to access a resource-A user is signed into the device and tries to access a protected resource on a Web server.

2. The system sends an HTTP or HTTPS GET request to the ACS-the system intercepts the request and checks whether it has already performed the necessary SSO operation to honor the request. If not, the system creates an authentication statement and passes an HTTP query variable called an artifact to the assertion consumer service.

An artifact profile is a Base64-encoded string that contains the source ID of the source site (that is, a 20-byte string that references the system) and a randomly generated string that acts as a handle to the authentication statement. (Note that a handle expires 5 minutes after the artifact is sent, so if the assertion consumer service responds after 5 minutes, the system does not send a statement. Also note that the system discards a handle after its first use to prevent the handle from being used twice.)

3. The ACS sends a SAML request to the system-The assertion consumer service uses the source ID sent in the previous step to determine the location of the device. Then the assertion consumer service sends a statement request wrapped in a SOAP message to the following address on the system:

`https://< <ivehostname>/danaws/saml.ws`

The request includes the statement handle passed in the previous step.



The system only supports type 0x0001 artifacts. This type of artifact passes a reference to the source site's location (that is, the source ID of the system), rather than sending the location itself. To handle type 0x0001 artifacts, the assertion consumer service must maintain a table that maps source IDs to the locations of partner source sites.

4. The system sends an authentication statement to the ACS-the system uses the statement handle in the request to find the correct statement in the system cache and then sends the appropriate authentication statement back to the assertion consumer service. The unsigned statement contains the user's identity and the mechanism he used to sign into the device.
5. The ACS sends a cookie to the system-The assertion consumer service accepts the statement and then it sends a cookie back to the system that enables the user's session.
6. The system sends the cookie to the Web server-the system caches the cookie to handle future requests. Then the system sends the cookie in an HTTP request to the Web server whose domain name matches the domain in the cookie. The Web server honors the session without prompting the user for credentials.



If you configure the system to use artifact profiles, you must install the Web server certificate on the assertion consumer service.

To write a SAML SSO artifact profile resource policy:

1. In the admin console, select **Users > Resource Policies > Web**.
2. If your administrator view is not already configured to show SAML policies, make the following modifications:
 1. Click the **Customize** button in the upper right corner of the page.
 2. Select the **SSO** check box.
 3. Select the **SAML** check box below the SSO check box.
 4. Click **OK**.
3. Use the tabs to display the **SSO > SAML** page.
4. Click **New Policy**.
5. On the New Policy page, enter:
 1. A name to label this policy.
 2. A description of the policy (optional).
6. In the Resources section, specify the resources to which this policy applies.
7. In the Roles section, specify:
 - **Policy applies to ALL roles**-To apply this policy to all users.
 - **Policy applies to SELECTED roles**-To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
 - **Policy applies to all roles OTHER THAN those selected below**-To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
8. In the Action section, specify:
 - **Use the SAML SSO defined below**-The system performs a single-sign on (SSO) request to the specified URL using the data specified in the SAML SSO details section. The system makes the SSO request when a user tries to access a SAML resource specified in the Resources list.
 - **Do NOT use SAML**-The system does not perform an SSO request.

- **Use Detailed Rules**-To specify one or more detailed rules for this policy.

9. In the SAML SSO Details section, specify:

- **SAML Assertion Consumer Service URL**-Enter the URL that the system should use to contact the assertion consumer service (that is, the access management server). For example, `https://<hostname>:<port>/dana-na/auth/saml-consumer.cgi`. (Note that the system also uses this field to determine the SAML recipient for its assertions.)



If you enter a URL that begins with HTTPS, you must install the assertion consumer service's root CA on the system.

- **Profile**-Select Artifact to indicate that the assertion consumer service should "pull" information from the system during SSO transactions.
- **Source ID**-Enter the source ID for the system. It must be a Base64-encoded string. The system decodes it and ensures that it is 20 bytes. You can use OpenSSL or other Base64 tool to generate the Base64-encoded string.



The system identifier (that is, the source ID) must map to the following URL on the assertion consumer service: `https://<ivehostname>/dana-ws/saml.ws`

- **Issuer**-Enter a unique string that the system can use to identify itself when it generates assertions (typically its hostname).



You must configure the assertion consumer service to recognize the unique string.

10. In the User Identity section, specify how the system and the assertion consumer service should identify the user:

- **Subject Name Type**-Specify which method the system and assertion consumer service should use to identify the user:
 - **DN**-Send the username in the format of a DN (distinguished name) attribute.
 - **Email Address**-Send the username in the format of an e-mail address.
 - **Windows**-Send the username in the format of a Windows domain qualified username.
 - **Other**-Send the username in another format agreed upon by the system and the assertion consumer service.

- **Subject Name**-Use variables to specify the username that the system should pass to the assertion consumer service. Or, enter static text.



You must send a username or attribute that the assertion consumer service will recognize.

11. In the Web Service Authentication section, specify the authentication method that the system should use to authenticate the assertion consumer service:
 - **None**-Do not authenticate the assertion consumer service.
 - **Username**-Authenticate the assertion consumer service using a username and password. Enter the username and password that the assertion consumer service must send.
 - **Certificate Attribute**-Authenticate the assertion consumer service using certificate attributes. Enter the attributes that the assertion consumer service must send (one attribute per line). For example, use cn=sales. You must use values that match the values contained in the assertion consumer service certificate.



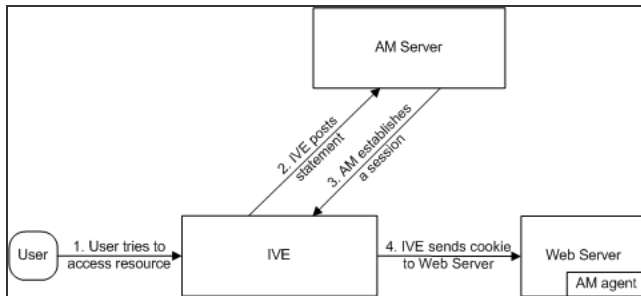
If you select this option, you must install the assertion consumer service root CA on the system.

12. **Cookie Domain**-Enter a comma-separated list of domains to which we send the SSO cookie.
13. Click **Save Changes**.
14. On the SAML SSO Policies page, order the policies according to how you want the system to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Creating a SAML 1.1 SSO POST Profile

When you choose to communicate using a POST profile (also called browser/POST profile), the system "pushes" authentication data to the access management system using an HTTP POST command over an SSL 3.0 connection.

The following figure shows the SAML communication process when the implementation uses the POST profile:



The system and an access management system use the following process to pass information:

1. The user tries to access a resource-A user is signed into the device and tries to access a protected resource on a Web server.
2. The system posts a statement-the system intercepts the request and checks whether it has already performed the necessary SSO operation to honor the request. If not, the system creates an authentication statement, digitally signs it, and posts it directly to the access management server. Since the statement is signed, the access management server must trust the certificate authority that was used to issue the certificate. Note that you must configure which certificate the system uses to sign the statement.
3. The AM establishes a session-If the user has the proper permissions, the access management server sends a cookie back to the system that enables the user's session.
4. The system sends the cookie to the Web server-the system caches the cookie to handle future requests. Then the system sends the cookie in an HTTP request to the Web server whose domain name matches the domain in the cookie. The Web server honors the session without prompting the user for credentials.



If you configure the system to use POST profiles, you must install the assertion consumer service's root CA on the system and determine which method the assertion consumer service uses to trust the certificate.

To write a SAML SSO POST profile resource policy:

1. In the admin console, select Users > Resource Policies > Web.
2. **If your administrator view is not already configured to show SAML policies, make the following modifications:**
 1. Click the Customize button in the upper right corner of the page.
 2. Select the SSO check box.

3. Select the SAML check box below the SSO check box.
4. Click **OK**.
5. Use the tabs to display the **SSO > SAML** page.
6. Click **New Policy**.
7. On the SAML SSO Policy page, enter:
 - A name to label this policy.
 - A description of the policy (optional).
8. In the Resources section, specify the resources to which this policy applies.
9. In the Roles section, specify:
 - Policy applies to ALL roles-To apply this policy to all users.
 - Policy applies to SELECTED roles-To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
 - Policy applies to all roles OTHER THAN those selected below-To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
10. In the Action section, specify:
 - Use the SAML SSO defined below-The system performs a single-sign on (SSO) request to the specified URL using the data specified in the SAML SSO details section. The system makes the SSO request when a user tries to access a SAML resource specified in the Resources list.
 - Do NOT use SAML-The system does not perform an SSO request.
 - Use Detailed Rules-To specify one or more detailed rules for this policy.
11. In the SAML SSO Details section, specify:
 - SAML Assertion Consumer Service URL-Enter the URL that the system should use to contact the assertion consumer service (that is, the access management server). For example, use https://hostname/acs.
 - Profile-Select POST to indicate that the system should "push" information to the assertion consumer service during SSO transactions.

- Issuer-Enter a unique string that the system can use to identify itself when it generates assertions. Typically, the issuer string is a hostname.



You must configure the assertion consumer service to recognize the unique string.

3. Signing Certificate-Specify which certificate the system should use to sign its assertions.

1. In the User Identity section, specify how the system and the assertion consumer service should identify the user:
2. Subject Name Type-Specify which method the system and assertion consumer service should use to identify the user:
 - DNDN-Send the username in the format of a DN (distinguished name) attribute.
 - Email Address-Send the username in the format of an e-mail address.
 - Windows-Send the username in the format of a Windows domain qualified username.
 - Other-Send the username in another format agreed upon by the system and the assertion consumer service.
3. Subject Name-Use variables to specify the username that the system should pass to the assertion consumer service. Or, enter static text.



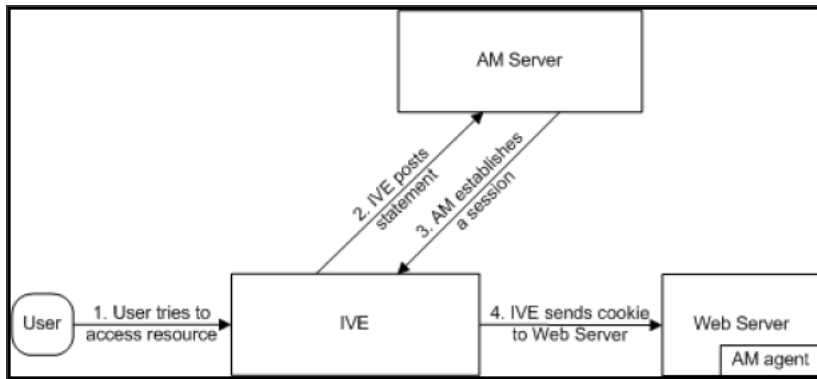
You must send a username or attribute that the assertion consumer service will recognize.

4. Cookie Domain-Enter a comma-separated list of domains to which we send the SSO cookie.
5. Click Save Changes.
6. On the SAML SSO Policies page, order the policies according to how you want the system to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Creating a SAML 1.1 ACL Resource Policy

When enabling access control transactions to a trusted access management system, the system and trusted access management system exchange information using the method shown in the following figure.

The following figure depicts the Access Control Policies:



The system and an access management system use the following process to pass information:

1. The user tries to access a resource-A user is signed into the system and tries to access a protected resource on a Web server.
2. The system posts an authorization decision query-If the system has already made an authorization request and it is still valid, the system uses that request. (The authorization request is valid for the period of time specified in the admin console.) If it does not have a valid authorization request, the system posts an authorization decision query to the access management system. The query contains the user's identity and the resource that the access management system needs to authorize.
3. The access management system posts an authorization decision statement-The access management system sends an HTTPS POST containing a SOAP message that contains the authorization decision statement. The authorization decision statement contains a result of permit, deny, or indeterminate.
4. The system sends the request to the Web browser-If the authorization decision statement returns a result of permit, the system allows the user access. If not, the system presents an error page to the user telling him that he does not have the proper access permissions.



If you configure the system to use access control transactions, you must install the SAML Web service root CA on the system.

To create a SAML access control resource policy:

1. In the admin console, select **Users > Resource Policies > Web**.
- 2.

If your administrator view is not already configured to show SAML policies, make the following modifications:

1. Click the **Customize button** in the upper right corner of the page.
2. Select the SAML ACL check box below the Access check box.
3. Click **OK**.
4. Use the tabs to display the **Access > SAML ACL** page.
5. On the SAML Access Control Policies page, click New Policy.
6. On the **New Policy** page, enter:
 - A name to label this policy.
 - A description of the policy (optional).
3. In the Resources section, specify the resources to which this policy applies.
4. In the Roles section, specify:
 - **Policy applies to ALL roles**-To apply this policy to all users.
 - **Policy applies to SELECTED roles**-To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
 - **Policy applies to all roles OTHER THAN those selected below**-To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
5. In the Action section, specify:
 - **Use the SAML Access Control checks defined below**-The system performs an access control check to the specified URL using the data specified in the SAML Access Control Details section.
 - **Do not use SAML Access**-The system does not perform an access control check.
 - **Use Detailed Rules**-To specify one or more detailed rules for this policy.
6. In the SAML Access Control Details section, specify:
 - **SAML Web Service URL**-Enter the URL of the access management system's SAML server. For example, use https://hostname/ws.

- **Issuer**-Enter the hostname of the issuer, which in most cases is the hostname of the access management system.



You must enter a unique string that the SAML Web service uses to identify itself in authorization assertions.

7. In the User Identity section, specify how the system and the SAML Web service should identify the user:

- **Subject Name Type**-Specify which method the system and SAML Web service should use to identify the user:
 - **DN**-Send the username in the format of a DN (distinguished name) attribute.
 - **Email Address**-Send the username in the format of an e-mail address.
 - **Windows**-Send the username in the format of a Windows domain qualified username.
 - **Other**-Send the username in another format agreed upon by the system and the SAML Web service.
 - **Subject Name**-Use variables to specify the username that the system should pass to the SAML Web service. Or, enter static text.



You must send a username or attribute that the SAML Web service will recognize.

8. In the Web Service Authentication section, specify the authentication method that the SAML Web service should use to authenticate the system:

- **None**-Do not authenticate the system.
- **Username**-Authenticate the system using a username and password. Enter the username and password that the system must send the Web service.
- **Certificate Attribute**-Authenticate the system using a certificate signed by a trusted certificate authority. If you have more than one certificate installed on the system, use the drop-down list to select which certificate to send to the Web service.



If you select this option, you must install the Web server certificate on the access management system Web server and determine which method the SAML Web service uses to trust the certificate.

9. In the Options section, specify:

- **Maximum Cache Time**-You can eliminate the overhead of generating an authorization decision each time the user requests the same URL by indicating that the system must cache the access management system's authorization responses. Enter the amount of time the system should cache the responses (in seconds).
- **Ignore Query Data**-By default, when a user requests a resource, the system sends the entire URL for that resource (including the query parameter) to the SAML Web service and caches the URL. You can specify that the system should remove the query string from the URL before requesting authorization or caching the authorization response.

10. Click **Save Changes**.

11. On the SAML Access Control Policies page, order the policies according to how you want the system to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.


Device Access Management Framework

The device access management framework leverages mobile device management (MDM) services so that you can use familiar Pulse client policies to enforce security objectives based on your device classification scheme: device status is MDM enrollment complete or incomplete; device status is MDM-policy compliant or non-compliant; device is employee owned or company owned; device platform is iOS, Android, or neither; and so forth.

In this framework, the MDM is a device authorization server, and MDM record attributes are the basis for granular access policy determinations. For example, you can implement policies that allow devices that have a clean MDM posture assessment and are compliant with MDM policies to access the network, but deny access to servers when you want to prevent downloads to employee-owned devices or to a particular platform that might be vulnerable. To do this, you use the device attributes and status maintained by the MDM in Pulse client role-mapping rules, and specify the device-attribute-based roles in familiar Pulse client policies.

The framework simply extends the user access management framework realm configuration to include use of device attributes as a factor in role mapping rules. The following figure illustrates the similarities.

The following figure depicts the User Access Management Framework and Device Access Management Framework:



For details about the deployment and configuration, refer to the *ICS Integration with MDM Servers Deployment Guide* available on the <https://www.ivanti.com/support/product-documentation> site.

Hosted Java Applets Templates

About Hosted Java Applet Templates

The Java applet upload feature enables you to store the Java applets of your choice directly on the device without employing a separate Web server to host them. When you use this feature, you simply upload the applets to the device (along with additional files that the applets reference) and create a simple Web page through the system that references the files. Then, the system intermediates the Web page and Java applet content using its Content Intermediation Engine.

For example, you might want to use the system to intermediate traffic between an IBM AS/400 system on your network and individual 5250 terminal emulators on your users' computers. To configure the system to intermediate this traffic, obtain the 5250 terminal emulator's Java applet. Then you can upload this applet to the system and create a simple Web page that references the applet. After you create the Web page through the system, it creates a corresponding bookmark that users can access through their home pages.

The system enables you to host Java applets using Web resource profile templates (described in these topics) as well as through Terminal Services resource profiles.

The hosted Java applets feature is a standard feature on all Ivanti Connect Secure devices.

Task Summary: Hosting Java Applets

The Java applet upload feature enables you to store the Java applets of your choice directly on the device without employing a separate Web server to host them.

To host Java applets on the device:

1. Specify which applets you want to upload, create bookmarks that reference the uploaded applets, and specify which roles can access the bookmarks using settings in the Users > Resource Profiles > Web page of the admin console.
2. (Optional.) To sign your Java applets, Select System > Configuration > Certificates > Code-Signing Certificates in the admin console to upload the Java certificate to the device. If you choose to skip this step, the user sees an untrusted certificate warning each time he accesses the corresponding bookmark.
3. (Optional.) To improve the performance of your Java applications:

- Select Enable Java instrumentation caching on the Maintenance > System > Options page of the admin console. This option can improve the performance of downloading Java applications.
- After you finish configuring the system, cache your Java applet and access it as an end user. This action eliminates the performance hit that occurs through the intermediation engine when the first end user accesses the applet.

Uploading Java Applets to Ivanti Connect Secure

You can use Java applets to intermediate traffic to various types of applications through the system. For example, you can upload the 3270 applet, 5250 applet, or Citrix Java applet. These applets enable users to establish sessions to IBM mainframes, AS/400s, and Citrix MetaFrame servers through terminal emulators. (Note that to enable the Citrix Java ICA client through a session, you must upload multiple Citrix .jar and .cab files to the device.)

The system enables you to upload individual .jar and .cab files or .zip, .cab, or .tar archive files. Archive files can contain Java applets and files referenced by the applets. Within the .zip, .cab, or .tar file, the Java applet must reside at the top level of the archive. You can upload any number of files to the system as long as their combined size does not exceed 100 MB.

To ensure compatibility with both Sun and Microsoft Java Virtual Machines (JVMs), you must upload both .jar and cab files to the device. (The Sun JVM uses .jar files, whereas the Microsoft JVM uses .cab files.)



When you upload Java applets, the system asks you to read a legal agreement before it finishes installing the applets. Read this agreement carefully-it obligates you to take full responsibility for the legality, operation, and support of the Java applets that you upload.

You can only upload 100 MB of Java applets to the system. The system displays the size of each applet that you upload on the Java Applets page, so you can determine, if necessary, which applets you want to delete.

Uploading Java applets requires signed ActiveX or signed Java applets to be enabled within the browser to download, install, and launch the client applications.

Signing Uploaded Java Applets

Unlike other Java applets that users can access through the system, you do not have to create a separate code-signing policy for the Java applets that you upload. The system automatically signs (or re-signs) them using the appropriate code-signing certificate. A code-signing certificate (also called an applet certificate) is a type of server-side certificate that re-signs Java applets intermediated by the system.

The system automatically signs (or resigns) your hosted Java applets with the code-signing certificate that you install through the System > Configuration > Certificates > Code-signing Certificates page of the admin console. If you do not install a code-signing certificate on the system, it uses its self-signed applet certificate to sign or re-sign the applets. In this case, users see an "untrusted certificate issuer" warning whenever they access the Java applets through the system.



The system re-instruments and re-signs your uploaded Java applets whenever you change (that is, import, renew, or delete) the corresponding code-signing certificate.

Creating HTML Pages That Reference Uploaded Java Applets

When uploading a Java applet to the system, you must create a simple Web page that references the applet. Users can access this Web page through a bookmark on their home pages or from external Web servers.

The Web page must contain a simple HTML page definition that references the uploaded Java applet. The Web page can also contain any additional HTML and JavaScript that you choose. The system can generate some of the Web page for you, including the HTML page definition and the references to your Java applet. (Note, however, that the system is not aware of all the applet-specific parameters that are required by your applet—you must find and fill these parameters in yourself.) When the system generates this HTML, it creates placeholders for any undefined values and prompts you to fill in the necessary values.

You can create these Web pages through Java applet upload resource profiles.

Accessing Java Applet Bookmarks

Users can access the applets you upload to the system using two methods:

- Bookmarks on the end-user console-When you create a Web page that references your uploaded Java applets, the system creates a corresponding link to the Web page and displays that link in the Bookmarks section of the end-user console. Users who map to the appropriate role can simply click the link to access the Java applet.
- Links on external Web servers-Users can link to the Java applet bookmarks from an external Web server by simply using the correct URLs. When the user enters a bookmark's URL (or clicks an external link that contains the URL), the system prompts the user to enter his username and password. If he properly authenticates, it allows him to access the bookmark. You can construct the URL to the Java applet bookmark using the syntax described in either of the following lines:

https://SecureAccessGateway_hostname/dana/home/launchwebapplet.cgi?

bmname=bookmark Name

https://SecureAccessGateway_hostname/dana/home/launchwebapplet.cgi?

id= <resourceID>&bmname=bookmarkName

You can determine the ID for a Java applet bookmark by accessing it through the home page and then extracting the ID from the Web browser's address bar.



Although the system enables you to create multiple bookmarks with the same name, we strongly recommend that you use a unique name for each. If multiple bookmarks have the same name and a user accesses one of these bookmarks using a URL that includes the bmname parameter, the system randomly picks which of the identically named bookmarks to display to the user. Also note that the bmname parameter is case-sensitive.

If you create links on external servers to Java applet bookmarks on the system and you are using multiple customized sign-in URLs, some restrictions occur.

Creating a Hosted Java Applet Resource Profile

To create a hosted Java applet resource profile:

1. Select **Users > Resource Profiles > Web in the admin console.**
2. Click **New Profile.**
3. Select **Hosted Java Applet** from the Type list.
4. Enter a unique name and optionally a description for the resource profile.

5. Select the Java applet that you want to associate with the resource profile from the Applet to use list. Or, if the applet that you want to use is not currently available in the list, click Edit Applet. Then:

1. Click **New Applet** to add an applet to this list. Or, select an existing applet and click Replace (to replace an existing applet with a new applet) or Delete (to remove an applet from the system.)



If you replace an existing archive, make sure that the new applet archive contains all of the necessary files for the applet to successfully launch and run. If the associated HTML for the applet refers to files that do not exist in the new archive, then the applet will not function correctly.

The system only allows you to delete applets that are not currently in use by a Web or Terminal Services resource profile.

2. Enter a name to identify the applet in the Name box (for new and replaced applets only).
3. Browse to the applet that you want to upload. You can upload applets (.jar or .cab files) or archives (.zip, .jar, and .tar files) that contain applets and all of the resources that the applets need (for new and replaced applets only).
4. Select the **Uncompress jar/cab file** check box if the file that you selected is an archive that contains the applet (New and replaced applets only).
5. Click OK and then click **Close Window**.



When you select an applet in the Java Applets dialog box, you are loading third-party software onto the Ivanti product. By clicking OK, you are agreeing to the following terms on behalf of yourself (as purchaser of the equipment) or the organization that purchased the Ivanti product, as applicable.

By loading third party software onto the Ivanti product, you are responsible for obtaining all rights necessary for using, copying, and/or distributing such software in or with the Ivanti product. Ivanti is not responsible for any liability arising from use of such third-party software and will not provide support for such software. The use of third-party software may interfere with the proper operation of the Ivanti product and/or Ivanti software, and may void any warranty for the Ivanti product and/or Ivanti software.

6. Use settings in the Autopolicy: Java Access Control section to enable access if your Java applets need to make socket connections.
7. Click **Save and Continue**.
8. Select the roles to which the resource profile applies In the Roles tab and click **Add**.

The selected roles inherit the autopolicies and bookmarks created by the resource profile. If it is not already enabled, the system also automatically enables the Web option in the Users > User Roles > Select_Role > General > Overview page of the admin console and the Allow Java Applets option Users > User Roles > Select_Role > Web > Options page of the admin console for all of the roles you select.

9. Click **Save Changes**.
10. Create bookmarks in the Bookmarks tab.

Configuring Hosted Java Applet Resource Profile Bookmarks

You must create bookmarks to your hosted Java applets to enable end users to access the applets.

To configure hosted Java applet resource profile bookmarks:

1. Select **Users > Resource Profiles > Web > Select Resource Profile > Bookmarks** in the admin console.
2. Click the appropriate link in the Bookmark column if you want to modify an existing bookmark. Or, click New Bookmark to create an additional bookmark.



Although it is generally easiest to create a resource profile session bookmark through the resource profile configuration page, you can choose to create one through the user roles page as well if you have already created a resource profile.

3. Enter a name and optionally a description for the bookmark. This information displays on the home page. (By default, the system names the bookmark the same name as the corresponding resource profile.)



We strongly recommend that you use a unique name for each bookmark to make it clear to users which link they are accessing.

4. Click Generate HTML to create an HTML page definition that includes references to your Java applets. Then, fill in any required attributes and parameters.

If you are using HTML generated by the system, make sure to search the HTML code for "`__PLEASE_SPECIFY__`" and update the code as necessary.

You can also add more HTML or JavaScript to this Web page definition. the system rewrites all of the code that you enter in this field



Make sure to enter unique HTML in this field. If you create two bookmarks with the same HTML code, the system deletes one of the bookmarks in the end-user view. You will still be able to see both bookmarks, however, in the administrator console.

5. List those attributes in the Multi-Valued User Attributes box if your HTML code contains attributes that may expand to multiple values (such as `userAttr.hostname` or `userAttr.ports`). When the user signs into a device, the system evaluates these attributes and creates separate bookmarks as necessary based on each of the individual values. If you use an attribute that expands to multiple values, but do not enter that attribute in this box, the system creates a single bookmark based on the attribute's first value.
6. Under Display options, click Bookmark opens new window to automatically open the Web resource in a new browser window. Note that this functionality applies only to role bookmarks and not bookmarks created by users. Next, select the following options if you want to hide UI elements from the user:
 - Do not display the browser address bar-Select this option to remove the address bar from the browser window. This feature forces all Web traffic through the system by precluding users in the specified role from typing a new URL in the address bar, which circumvents the system.
 - Do not display the browser toolbar-Select this option to remove the menu and toolbar from the browser. This feature removes all menus, browsing buttons, and bookmarks from the browser window so that the user browses only through the system.

7. Under **Roles**, specify the roles to which you want to display the bookmark if you are configuring the bookmark through the resource profile pages:
 - **ALL selected roles**-Select this option to display the bookmark to all of the roles associated with the resource profile.
 - **Subset of selected roles**-Select this option to display the bookmark to a subset of the roles associated with the resource profile. Then select roles from the **ALL Selected Roles** list and click **Add** to move them to the **Subset of selected roles** list.
8. Click **Save Changes**.

Creating Hosted Java Applets Bookmarks Through the User Roles Page

It is generally easiest to create a hosted Java applets bookmark through the resource profile configuration pages, as explained in previous topic. However, you can choose to create a resource profile session bookmark through the user roles page using the following instructions:

1. Select **Users > Roles > Select_Role > Web > Bookmarks in the admin console**.
2. Click **New Bookmark**.
3. Select **Pick a Web Resource Profile** from the **Type** list. (The system does not display this option if you have not already created a hosted Java applet resource profile.)
4. Select an existing resource profile.
5. Click **OK**. (If you have not already associated the selected role with the resource profile, the system automatically makes the association for you. The system also enables any access control policies for the role that are required by the resource profile.)
6. If this role is not already associated with the selected resource profile, the system displays an informational message. If you see this message, click **Save Changes** to add this role to the resource profile's list of roles and to update the profile's autopolicies as required. Then, repeat the previous steps to create the bookmark.
7. Configure the bookmark settings.



When you create a resource profile bookmark through the user roles page (instead of the standard resource profiles page), the system only associates the generated bookmark with the selected role. The system does not assign the bookmark to all of the roles associated with the selected resource profile.

Required Attributes for Uploaded Java Applets

When you create a Java applets bookmark through the system, you must define the following attributes and their corresponding values. If you use the Generate HTML feature, it populates some of this information for you and adds PLEASE_SPECIFY to those attributes whose values you must specify. When specifying attributes and their corresponding values, use the attribute="value" format.



The system generates parameters that it knows are required. Note, however, that it is not aware of all the applet-specific parameters that are required by your applet—you must find and fill in these parameters yourself.

Attributes that are required by the system include:

- **codebase**—Indicates where the Web browser can fetch the applet. Use the <<CODEBASE>> variable, which points to the location on the system where it stores the Java applet. When entering a path to a file, note that <<CODEBASE>> includes a trailing slash, which means the following example works:

```

```

This example does not work:

```

```

- **archive**—Indicates which archive file (that is, .jar, .cab, or .zip file) the Web browser should fetch. Example:

```
archive="JICAEngN.jar"
```

In addition to the required attributes listed earlier, you may also use the following optional attributes when creating a Java applet bookmark:

- **name**—Specifies a label for the Java applet. Example:

```
name="CitrixJICA"
```

- **host**—Specifies, for terminal sessions, the server to which the system should connect.
- **port**—Specifies, for terminal sessions, the port to which the system should connect.
- **width and height**—Indicates the size of the Java applet window. Example:

```
width="640" height="480"
```

- **align**—Indicates the Java applet window's alignment within the browser window. Example:

```
align="top"
```



When defining attributes and their corresponding values, note the following:

- We strongly recommend that you not include `useslibrarycabbase` parameter in the HTML, because it causes the cab file to be permanently installed on the user's machine. If you later change a cab file on the system, all users will have to manually delete the cab files on their machines to get the new version from the system.
- We do not support applet tags that are constructed through the `document.write` function because the dynamic HTML interferes with the system parser.
- We do not support relative links to URLs, documents, or images in your HTML. If you do, the links will break when the user tries to access them from the end-user console. Instead, you should include absolute links. If you are linking to a document or image included in your zip file, use the `<<CODEBASE>>` variable to indicate that the system can find the file in the uploaded zip archive. For example:

```

```

Required Parameters for Uploaded Java Applets

When you create a Java applets bookmark through the system, you must specify parameters and values that should be passed to the Java applet. These parameters are completely applet-specific. When specifying parameters and their corresponding values, use the following format:

```
<param name="parameterName" value="valueName">
```

Where all of the text is literal except `parameterName` and `valueName`.

You can use variables to pass values to the Java applet by enclosing the variable names in double-brackets. For example, you might choose to pass the `<<username>>` and `<<password>>` values to the Java applet.



When using the Java applet upload feature, if you include the `<password>` token within the generated HTML, it appears as cleartext if you view the source in the browser window that launches the applet. This behavior cannot be changed because the system does not control how the Java applet processes the password. We strongly discourage the use of the `<password>` token in the HTML code.

If you find a Web page that contains an applet that you want to use, go to the demonstration site and view the source on the page that runs the Java applet. Within the source, look at the applet tag. Pick out the code attribute in the source and determine if it contains any special parameters that you need to pass to the browser. In most cases, you should be able to copy and paste the code attribute and its corresponding parameters directly into the HTML field for your bookmark. Note, however, that if a parameter references a resource on the local Web server, you cannot copy and paste the reference into the bookmark because the system does not have access to the other Web server's local resources. When copying and pasting parameters from another source, always check the values of the parameters.

Resource Policies

Resource Policies

A resource policy is a system rule that specifies resources and actions for a particular access feature. A resource is either a server or file that can be accessed through the system, and an action is to "allow" or "deny" a resource or to perform or not perform a function. Each access feature has one or more types of policies, which determine the system's response to a user request or how to enable an access feature. You may also define detailed rules for a resource policy, which enable you to evaluate additional requirements for specific user requests.

You can create the following types of resource policies through the Resource Policies pages:

- **Web Resource Policies** - The Web resource policies specify the Web resources to which users may or may not browse. They also contain additional specifications such as header caching requirements, servers to which java applets can connect, code-signing certificates that the system should use to sign java applets, resources that the system should and should not rewrite, applications for which the system performs minimal intermediation, and single sign-on options.



From 9.1R3 release, for a fresh installation, the predefined Web Access Resource Policy "Initial Policy for Local Resources" is in "Deny" state by default.

-
- **File Resource Policies** - The file resource policies specify the Windows, UNIX, and NFS file resources to which users may or may not browse. They also contain additional specifications such as file resources for which users need to provide additional credentials.



From 9.1R3 release, for a fresh installation, the predefined Windows File Access Resource Policy "Initial File Browsing Policy" is in "Deny" state by default.

-
- **Secure Application Manager Resource Policies** - The Secure Application Manager resource policies allow or deny access to applications configured to use JSAM or PSAM to make socket connections.
 - **Terminal Services Policies** - The Terminal Services resource policies allow or deny access to the specified Windows servers or Citrix Metaframe servers.
 - **VPN Tunneling Resource Policies** - The VPN Tunneling resource policies allow or deny access to the specified servers and specify IP address pools.



You can also create resource policies as part of the resource profile configuration process. In this case, the resource policies are called "advanced policies."

Resource policies are an integral part of the access management framework, and therefore are available on all Ivanti Connect Secure products. However, you can access only resource policy types that correspond to your licensed features.

Resource Policy Components

A resource policy contains the following information:

- **Resources** - A collection of resource names (URLs, hostnames, or IP address/netmask combinations) that specifies the resources to which the policy applies. You can specify a resource using a wildcard prefix to match hostnames. The default resource for a policy is star (*), meaning that the policy applies to all related resources.
- **Roles** - An optional list of user roles to which this policy applies. The default setting is to apply the policy to all roles.
- **Action** - The action for the system to take when a user requests the resource corresponding to the Resource list. An action may specify to allow or deny a resource or to perform or not perform an action, such as to rewrite Web content or allow Java socket connections.
- **Detailed Rules** - An optional list of elements that specifies resource details (such as a specific URL, directory path, file, or file type) to which you want to apply a different action or for which you want to evaluate conditions before applying the action. You can define one or more rules and specify the order in which the system evaluates them.

Specifying Resources for a Resource Policy

The system platform's engine that evaluates resource policies requires that the resources listed in a policy's Resources list follow a canonical format. This section describes the canonical formats available for specifying Web, file, and server resources. When a user tries to access a specific resource, the system compares the requested resource to the resources specified in the corresponding policies, starting with the first policy in a policy list. When the engine matches a requested resource to a resource specified in a policy's Resources list, it then evaluates further policy constraints and returns the appropriate action to the appliance (no further policies are evaluated). If no policy applies, then the appliance evaluates the auto-allow bookmarks (if defined); otherwise the default action for the policy is returned.



You may not see the auto-allow option if you are using a new installation, if you use resource profiles rather than resource policies, or if an administrator has hidden the option.

General Notes About the Canonical Formats

Please note the following when using canonical formats:

- If a path component ends with forward-slash_star (/*), then it matches the leaf node and everything below. If the path component ends with forward-slash_percent (/%), then it matches the leaf node and everything one-level below only. For example:

```

/intranet/*matches:
/intranet
/intranet/home.html
/intranet/elee/public/index.html
/intranet/% matches:
/intranet
/intranet/home.html
but NOT /intranet/elee/public/index.html

```

- A resource's hostname and IP address are passed to the policy engine at the same time. If a server in a policy's Resources list is specified as an IP address, then the evaluation is based on the IP address. Otherwise, the engine tries to match the two hostnames. It does not perform a reverse-DNS-lookup to determine the IP.



You cannot specify a hostname for a VPN Tunneling resource policy. You can only specify an IP address.

- If a hostname is not fully qualified in the hosts file, such as "pulsesecure" instead of "intranet.pulsesecure.net", and you are accessing a hostname using the short name, then the engine performs the resource matching against the short name. If, however, the short name is not in the hosts file and the hostname resolution is done by DNS (by adding the domains listed in the Networks configuration page), then the fully qualified domain name (FQDN) is used for resource matching. In other words, for web resource policies a DNS lookup of the short name is performed. The result of the DNS lookup is a FQDN; the engine matches the FQDN with the ones entered in the UI.

Specifying Server Resources

When specifying server resources for Terminal Services, or VPN Tunneling resource policies, note the following guidelines.

The canonical format is **[protocol://] host [:ports]**

The components are:

- Protocol (optional) - Possible case-insensitive values:
 - tcp
 - udp
 - icmp

If the protocol is missing, then all protocols are assumed. If a protocol is specified, then the delimiter "://" is required. No special characters are allowed.



Available only to VPN Tunneling policies. For other access feature resource policies, such as Secure Application Manager, it is invalid to specify this component.

- Host (required) - Possible values:
 - IP address/Netmask - The IP address needs to be in the format: a.b.c.d
The netmask may be in one of two formats:
 - Prefix: High order bits
 - IP: a.b.c.d
 For example: 10.11.149.2/24 or 10.11.149.2/255.255.255.0
No special characters are allowed.
 - DNS Hostname - For example: www.pulsesecure.net

The following table lists the DNS Hostname Special Characters

*	Matches ALL characters
%	Matches any character except dot (.)
?	Matches exactly one character



You cannot specify a hostname for a VPN Tunneling resource policy. You can only specify an IP address.

- Ports (optional) - Possible values are shown in the following table:

*	Matches ALL ports; no other special characters are allowed
---	--

port[,port]*	A comma-delimited list of single ports. Valid port numbers are [1- 65535]. Do not enter a space between port numbers. You can specify up to 15 ports.
[port1]-[port2]	A range of ports, from port1 to port2, inclusive.



You may mix port lists and port ranges, such as: 80,443,8080-8090, except for in VPN Tunneling where mixing of port lists and port ranges is not supported.

If the port is missing, then the default port 80 is assigned for http, 443 for https. For VPN Tunneling, if the port is missing then the default port http is *. If a port is specified, then the delimiter ":" is required. For example:


```
<username>.danastreet.net:5901-5910
tcp://10.10.149.149:22,23
tcp://10.11.0.10:80
udp://10.11.0.10:*
```

Resource Policy Evaluation

When the system receives a user request, it evaluates the resource policies corresponding to the type of request. When it processes the policy that corresponds to the requested resource, it applies the specified action to the request. This action is defined on the policy's General tab or Detailed Rules tab. For example, if a user requests a Web page, the system knows to use the Web resource policies. In the case of Web requests, the system always starts with the Web Rewriting policies (Selective Rewriting and Pass through Proxy) to determine whether or not to handle the request. If none of these policies applies (or none is defined), the system then evaluates the Web Access policies until it finds one that pertains to the requested resource.

The system evaluates a set of resource policies for an access feature from the top down, meaning that it starts with the policy numbered one and then continues down the policy list until it finds a matching policy. If you defined detailed rules for the matching policy, the system evaluates the rules from the top down, starting with the rule numbered one and stopping when it finds a matching resource in the rule's Resource list. The following figure illustrates the general steps of policy evaluation:

The following figure depicts the Resource Policy Evaluation Steps:



Details regarding each evaluation step:

1. The system receives a user request and evaluates the user's session role to determine if the corresponding access feature is enabled. A user's "session role" is based on either the role or roles to which the user is assigned during the authentication process. The access features enabled for a user are determined by an authentication realm's role mapping configuration.
2. The system determines which policies match the request. The system evaluates the resource policies related to the user request, sequentially processing each policy until finding the one whose resource list and designated roles match the request. (If you configure the system using resource profiles, the system evaluates the advanced policies that you configure as part of the resource profile.)

The Web and file access features have more than one type of policy, so the system first determines the type of request (such as to a Web page, Java applet, or UNIX file) and then evaluates the policies related to the request. In the case of the Web access feature, the Rewriting policies are evaluated first for every Web request. The remaining access features - Secure Application Manager, Secure Terminal Access-have only one resource policy.

3. The system evaluates and executes the rules specified in the matching policies. You can configure policy rules to do two things:
 - Specify resources to which an action applies at a more granular level. For example, if you specify a Web server in the main policy settings for a Web Access resource policy, you can define a detailed rule that specifies a particular path on this server and then change the action for this path.
 - Require the user to meet specific conditions written as boolean expressions or custom expressions in order to apply the action.
4. The system stops processing resource policies as soon as the requested resource is found in a policy's Resource list or detailed rule.



If you use automatic (time-based) dynamic policy evaluation or you perform a manual policy evaluation, the system repeats the resource evaluation process described in this section.

Creating Detailed Rules for Resource Policies

The Web, file, Secure Application Manager, and VPN Tunneling access features enable you to specify resource policies for individual Web, file, and application servers. For these access features, you can specify any number of resource policies, and for each, you can define one or more detailed rules.

A detailed rule is an extension of a resource policy that may specify:

- Additional resource information - such as a specific path, directory, file, or file type - for resources listed on the General tab. Note that you may also specify the same resource list (as on the General tab) for a detailed rule if the only purpose of the detailed rule is to apply conditions to a user request.
- An action different from that specified on the General tab (although the options are the same).
- Conditions that must be true in order for the detailed rule to apply.

In many cases, the base resource policy - that is, the information specified on the General tab of a resource policy - provides sufficient access control for a resource:

If a user belonging to the (defined_roles) tries to access the (defined_resources), DO the specified (resource_action).

You may want to define one or more detailed rules for a policy when you want perform an action based on a combination of other information, which can include:

- A resource's properties - such as its header, content-type, or file type.
- A user's properties - such as the user's username and roles to which the user maps.
- A session's properties - such as the user's source IP or browser type, whether the user is running Host Checker, the time of day, and certificate attributes.

Detailed rules add flexibility to resource access control by enabling you to leverage existing resource and permission information to specify different requirements for different users to whom the base resource policy applies.

Writing a Detailed Rule for Resource Policies

Detailed rules add flexibility to resource access control by enabling you to leverage existing resource and permission information to specify different requirements for different users to whom the base resource policy applies.

To write a detailed rule for a resource policy:

1. On the New Policy page for a resource policy, enter the required resource and role information.
2. In the Action section, select **Use Detailed Rules** and then click **Save Changes**.
3. On the Detailed Rules tab, click **New Rule**.

4. On the Detailed Rule page:

1. In the Action section, specify:
 - **Disable SSO** - The system disables automatic SSO authentication for this user role and, instead, prompts the user for sign-in credentials.

-
- **Basic Auth** - This option specifies that the system use the Basic Authentication Intermediation method to control SSO behavior.
2. **Enable Intermediation** - Select the credentials to use. If this pull-down menu is blank, no basic authentication SSO settings are defined in the SSO General tab.
 3. **Disable Intermediation** - When you select this option, the system does not intermedate the challenge/response sequence.



The system always intermediates requests to Web proxies that require basic authentication, even if you select Disable Intermediation.

Although you are given an option to disable basic authentication intermediation, we do not recommend this option, as it is a very insecure authentication method and, in some cases, can transmit user credentials over the network in clear (unencrypted) text.

4. **NTLM** - This option specifies that the system use the Microsoft NTLM Intermediation method to control SSO behavior.

Select the credentials to use. If this pull-down menu is blank, no NTLM SSO settings are defined in the SSO General tab.

Select the **Fallback to NTLM V1** option to try both NTLM V1 and NTLM V2. If you do not select this option, the system falls back only to NTLM V2. An intermediation page appear if SSO fails.

5. **Kerberos** - This option specifies that the system use the Kerberos Intermediation method to control SSO behavior.

Select the credentials to use. If this pull-down menu is blank, no kerberos SSO settings are defined in the SSO General tab.

Select the **Fallback to NTLM V2** option to fallback only to NTLM V2 if kerberos fails. If you do not select this option, a Kerberos intermediation page appears if Kerberos SSO fails.

6. **Constrained Delegation** - This option specifies that the system use the constrained delegation intermediation method to control SSO behavior.

Select the credentials to use. If this pull-down menu is blank, no constrained delegation SSO settings are defined in the SSO General tab.

Select the **Fallback to Kerberos** option fallback to Kerberos if constrained delegation fails. If you select this option, an intermediation page appears if constrained delegation fails. If you do not select this option and constrained delegation fails, an error page appears.

7. In the Resources section, specify any of the following (required):
 - The same or a partial list of the resources specified on the General tab.
 - A specific path or file on the server(s) specified on the General tab, using wildcards when appropriate. For information about how to use wildcards within a Resources list, see the documentation for the corresponding resource policy.
 - A file type, preceded by a path if appropriate or just specify `/*.file_extension` to indicate files with the specified extension within any path on the server(s) specified on the General tab.
8. In the Conditions section, specify one or more expressions to evaluate in order to perform the action (optional):
 - Boolean expressions: Using system variables, write one or more boolean expressions using the NOT, OR, or AND operators.
 - Custom expressions: Using the custom expression syntax, write one or more custom expressions.



You can use the `<USER>` substitution variable in ACLs for web pages, files, and SAM. You cannot use the variable in VPN Tunneling ACLs.

When specifying a time condition, the specified time range cannot cross midnight. The workaround is to break the time range into two conditions.

9. Click **Save Changes**.
 5. On the **Detailed Rules** tab, order the rules according to how you want the system to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a rule's Resource list, it performs the specified action and stops processing rules (and other resource policies).

Customizing Resource Policy UI Views

You can limit which resource policies the system displays on any given resource policy page based on user roles. For instance, you can configure the Users > Resource Policies > Web page of the admin console to only display those resource policies that are assigned to the "Sales" user role.

To control which resource policies the system displays:

1. Navigate to **Users > Resource Policies > Policy Type**.
2. From the Show all policies that apply to list, select **All Roles** or an individual role.
3. Click **Update**. The system displays resource policies that are assigned to the selected roles.

Citrix Templates

About Citrix Templates

The system supports several mechanisms for intermediating traffic between a Citrix server and client, including the Citrix Services Client proxy, JSAM, PSAM, VPN Tunneling, and the hosted Java applets feature.

The Citrix Web template enables you to easily configure access to a Citrix server using the Citrix Services Client proxy, JSAM, or PSAM. The Citrix Web template is a resource profile that controls access to Citrix applications and configures Citrix settings as necessary. Citrix Web templates significantly reduce your configuration time by consolidating configuration settings into one place and by prepopulating a variety of resource policy settings for you depending on the type of Citrix setup you select.

Because of their highly simplified configurations, templates are the ideal Citrix configuration method if you want to deliver ActiveX or Java applets from a third-party Web server through the system.

Citrix Web templates simplify your configuration by automatically detecting whether the Citrix Web client or the Citrix Java applet is being used and employing the appropriate access mechanism accordingly.

We strongly recommend using Citrix templates instead of the traditional role and resource policy configuration options available through the system.



Ivanti does not support saving a Citrix application shortcut to the desktop through the system when the loopback IP address is running on the client. Double-clicking this shortcut returns an error as it does not use JSAM or PSAM.

Comparing Access Mechanisms for Configuring Citrix

Ivanti Connect Secure supports several mechanisms for intermediating traffic between a Citrix server and client, including the Citrix Terminal Services proxy, JSAM, PSAM, VPN Tunneling, and the hosted Java applets feature.

The following table describes Accessing Citrix Metaframe Server:

Requirement	Terminal Services	JSAM	PSAM
User experience	The user launches the published application by clicking the bookmark or icon in the Terminal Services section of the end user console.	JSAM auto-launches when the user signs into the device or the user launches JSAM manually. The user launches the published application using standard methods such as the Windows Start menu or a desktop icon.	PSAM auto-launches when the user signs into the device or the user launches PSAM manually. The user launches the published application using standard methods such as the Windows Start menu or a desktop icon.
Accessing published applications from Mac or Linux	Macintosh and Linux users cannot access published applications from a Citrix Metaframe server.	Macintosh and Linux users can access published applications from a Citrix Metaframe server.	Macintosh and Linux users cannot access published applications from a Citrix Metaframe server.

Requirement	Terminal Services	JSAM	PSAM
Admin configuration	You can specify which ports the system intermediates. If you do not configure this information, the system automatically monitors ports 1494 and 2598.	You cannot configure Citrix as a standard application. Instead, you need to create a custom JSAM application, provide the server names of all Metaframe servers, and specify which ports to monitor. This enables you to use applications such as Citrix Secure Gateways (CSGs) and published applications that use ports other than 1494.	You must specify which ports and applications the system monitors. This enables you to use applications such as Citrix Secure Gateways (CSGs) and published applications that use ports other than 1494.
Administrator privileges	If a Citrix Web client is not installed on the user's desktop, administrator privileges are required. This is a limitation of the installation of the Citrix client. To install and run the Citrix Services Client proxy, administrator privileges are not required.	Requires administrator privileges to run JSAM because etc/hosts file modifications are required.	Requires administrator privileges to install PSAM.
Modifying host file	Does not require modification of the etc/hosts file.	Requires modification of the etc/hosts file.	Does not require modification of the etc/hosts file.

Creating Resource Profiles for Citrix Storefront Server

If you have the Citrix StoreFront 3.1 and above, you can create a Web template to allow users to access Citrix applications without the need for a Citrix client. Users must have one of the following browser versions (or later) to support HTML5 and Websockets:

- Internet Explorer 10
- Safari 6
- Google Chrome 23
- Mozilla Firefox 17



You can collect all the logs related to this feature using hprewrite-server as the process name.

To create a resource profile using the Citrix template:

1. Select **Users > Resource Profiles > Web** in the admin console.
2. Click **New Profile**.
3. Select **Citrix StoreFront 3.1 and above** from the Type list.
4. Enter a unique name and optionally a description for the Citrix resource profile.
5. Enter the URL of the Citrix StoreFront Web server in the Base URL field. Use the format: [protocol://]host[:port][path]. The system uses the specified URL to define the default bookmark for the Citrix resource profile. You may enter a directory URL or a file URL.
6. Under **Citrix Settings**, select the **ICA Client Access** option. Admin can either choose to go with the HTML5 way of delivery or can choose to deliver ICA over CTS/PSAM/HTML5 Access clients. If admin chooses the ICA over CTS/PSAM/HTML5 Access, the corresponding ACL should be created and when ICS rewrites ICA content it should launch the appropriate client. Add the **Number of servers/applications** and **Citrix Ports** which require ICA client access.
7. Select the **Autopolicy: Web Access Control** check box to create a policy that allows or denies users access to a specific resource under the Base URL. Enter the full URL of the resource, select **Allow** or **Deny**, and click **Add**. By default, the system automatically creates a policy that enables access to the resource and all of its subdirectories.
8. Select the **Autopolicy: Single Sign-on** check box to automatically pass data such as usernames and passwords to the Citrix application. The system automatically adds the most commonly used values to the single sign-on autopolicy.

9. If you want to perform a form POST when a user makes a request to the resource specified in the Resource field, select the **POST the following data** check box and specify the following:
- In the Resource field, specify the application's sign-in page, such as: `http://my.domain.com/public/login.cgi`. Wildcard characters are not supported in this field. To automatically post values to a specific URL when an end user clicks on a system bookmark, the resource that you enter here must exactly match the URL that you specify in the Base URL field.
 - In the Post URL field, specify the absolute URL where the application posts the user's credentials, such as: `http://yourcompany.com/login.cgi`. You can determine the appropriate URL using a TCP dump or by viewing the application's sign-in page source and searching for the POST parameter in the FORM tag.
 - Select the **Deny direct login for this resource** check box if you do not want to allow users to manually enter their credentials in a sign-in page. Users may see a sign-in page if the form POST fails.)
 - Select the **Allow multiple POSTs to this resource** check box if you want to send POST and cookie values to the resource multiple times if required. If you do not select this option, the system does not attempt single sign-on when a user requests the same resource more than once during the same session.
 - Optionally specify the following for each item of user data you want to post and click **Add**:
 - **Label**-The name used to identify the data.
 - **Name**-The name used to identify the data in the Value field. The back-end application should expect this name.
 - **Value**-The value to post to the form for the specified Name. You can enter static data, a system variable, or system session variables containing username and password values.
 - **User modifiable?**-Select **Not modifiable** to prevent users from changing the information in the Value field. Select **User CAN change value to allow** users to specify data for a back-end application. Select **User MUST change value** if users must enter additional data to access a back-end application. If users can or must change the value, a field for data entry appears on the user's Advanced Preferences page. This field is labeled using the name in the Label field. If you enter a value in the Value field, this data appears in the field but is editable.

10. To post header data to the specified URL when a user makes a request to a resource specified in the Resource field, select the Send the following data as request headers check box. Then:

1. In the Resource section, specify the resources to which this policy applies.
2. Optionally specify the header data to post by entering data in the following fields and clicking **Add**:
 - **Header name**-The text to send as header data.
 - **Value**-The value for the specified header.
3. Click **Save** and **Continue**.
4. Select the roles in the Roles tab to which the Citrix resource profile applies and click **Add**.

The selected roles inherit the autopolicies and bookmarks created by the Citrix resource profile. If it is not already enabled, the system also automatically enables the Web option in the Users > User Roles > Select_Role > General > Overview page of the admin console and the Allow Java Applets option in the Users > User Roles > Select_Role > Web > Options page of the admin console for all of the roles you select.

5. Click **Save Changes**.
6. (Optional.) Select the **Bookmarks** tab to modify the default bookmark created by the system and/ or create new bookmarks. By default, the system creates a bookmark for the URL defined in the Base URL field and displays it to all users assigned to the role specified in the Roles tab.

Lotus iNotes Templates

Creating Resource Profiles Using the Lotus iNotes Template

A Lotus iNotes template is a resource profile that controls access to the Web application and configures iNotes settings as necessary. Lotus iNotes templates significantly reduce your configuration time by consolidating settings into one place and by prepopulating a variety of resource policy settings for you depending on the type of setup you select.

The system supports intermediating traffic to Lotus iNotes through a Web rewriting resource profile template, JSAM, PSAM, and VPN Tunneling. This topic describes how to configure access using the Web rewriting template. The prepopulated values vary depending on the version of iNotes you select and are based on the most common deployment of the servers.

To create a resource profile using the Lotus iNotes template:

1. Select **Users > Resource Profiles > Web** in the admin console.
2. Click **New Profile**.
3. Select the Lotus Notes version from the Type list.

The following figure depicts Creating Resource Profiles Using the Lotus iNotes Template:

The screenshot shows the 'New Web Application Resource Profile' configuration page. The 'Type' dropdown is set to 'Custom'. The 'Name' field contains 'Custom'. The 'Description' field is empty. The 'Base URL' field is empty. The 'Autopolicies' dropdown is open, showing options: 'Microsoft OWA 2016 and above', 'Lotus iNotes 7', 'Lotus iNotes 8', 'Lotus iNotes 8.5 and above', and 'Microsoft Sharepoint'. A red box highlights the 'Lotus iNotes 8' option. A blue button labeled 'Show ALL autopolicy types >>' is visible at the bottom of the dropdown menu.

4. Enter a unique name and optionally a description for the Lotus Notes resource profile.
5. Enter the URL of the Lotus iNotes resource to which you want to control access in the Base URL box. Use the format: [protocol://]host[:port][/path]. The system uses the specified URL to define the default bookmark for the Lotus iNotes resource profile. You may enter a directory URL or a file URL.

6. Under iNotes setting, select **Allow caching on client** to let Web browsers store non-user data, such as Javascript and CSS files, on a user's machine. Select **Minimize caching on client to allow** the system to send a cache-control:no-store header or a cache-control:no-cache header based on the user's Web browser and content type. This is the same as smart caching.

The Allow caching on client option caches content that the backend iNotes server typically caches. This caching option improves performance by using the cached content instead of retrieving the content from the server the next time the page displays. The Minimize caching on client option provides security by sending a cache-control:no-store header or a cache-control:no-cache header to either not store content or to re-validate the cached content each time it is requested. With both caching option, you can choose to either allow or prevent the uploading or downloading of attachments.

7. Select the **Prevent download of attachments** check box to prohibit users from downloading attachments to their systems. Select the Prevent upload of attachments check box to prevent users from transmitting (uploading) attachments to the system.
8. Select the **Autopolicy: Web Access Control** check box to create a policy that allows or denies users access to the Web resource (and all of its subdirectories) listed in the Resource field.
 1. In the Resource box, specify the Web server or HTML page to which you want to control access using the format: [protocol://]host[:port][]/path].
 2. From the Action list, select **Allow** to enable access to the specified resource or Deny to block access to the specified resource.
 3. Click **Add**.
9. Select the **Autopolicy: Caching** check box to specify the resources to which this policy applies in the Resource box.



The correct caching resource policy must be configured to allow end users to open and save e-mail attachments of different document types in iNotes. For example, if the caching policy is set to Smart, end users cannot save .htm or .html attachments to disk.

10. Select the **Autopolicy: Web Compression** check box to create a policy that specify which types of Web data the system should and should not compress.

1. In the Resources field, specify the resources to which this policy applies.
 2. Select one of the following options from the Action list:
 - Compress-Compresses the supported content types from the specified resource.
 - Do not compress-Do not compress the supported content types from the specified resource.
 3. Click **Add**.
11. Select the **Autopolicy: Single Sign-On** check box to pass data such as the username and password to the Lotus iNotes application.
1. Click **Save and Continue**.
 2. Select the roles to which the Lotus iNotes resource profile applies in the Roles tab and click **Add**.

The selected roles inherit the autopolicies and bookmarks created by the Lotus iNotes resource profile. If it is not already enabled, the system also automatically enables the Web option in the Users > User Roles > Select Role > General > Overview page of the admin console.
12. Click **Save Changes**.
13. (Optional.) In the Bookmarks tab, modify the default bookmark created by the system and/or create new ones.

Microsoft OWA Templates

Creating Resource Profiles Using the Microsoft OWA Template

A Microsoft Outlook Web Access (OWA) template is a resource profile that controls access to the application and configures OWA settings as necessary. OWA templates significantly reduce your configuration time by consolidating configuration settings into one place and by prepopulating a variety of resource policy settings for you depending on the type of setup you select.

The system supports intermediating traffic to Microsoft OWA through a Web rewriting resource profile template, JSAM, PSAM, and VPN Tunneling. This topic describes how to configure access using the Web rewriting template. The prepopulated values vary depending on the version of OWA you select and are based on the most common deployment of the servers.

The following figure depicts Creating Resource Profiles Using the Microsoft OWA Template (**Users > Web Application Resource Profiles > New Web Application Resource Profiles**):

To create a resource profile using the Microsoft OWA template:

1. Select **Users > Resource Profiles > Web Applications/Pages** in the admin console.
2. Click **New Profile**.
3. Select your Microsoft OWA version from the Type list.
4. Enter a unique name and optionally a description for the Citrix resource profile.
5. Enter the URL of the OWA resource to which you want to control access In the Base URL box. Use the format: [protocol://]host[:port][path]. The system uses the specified URL to define the default bookmark for the OWA resource profile. You may enter a directory URL or a file URL.
6. Under OWA settings select the following options,

- (OWA 2010,2013, and 2016 and above.) Select **Managed Device** to cache files. If you configure a Form post SSO, the trusted parameter is set to 4. This indicates the end user's device is private.
- (OWA 2010,2013, and 2016 and above.) Select **Unmanaged Device** to not cache files. If you configure a Form post SSO, the trusted parameter is set to 0. This indicates the end user's device is public.



If it is necessary to download an attachment, the file is cached even though you select Unmanaged Device.

- Select **Prevent download of attachments** to prohibit users from downloading attachments to their systems.
 - Select **Prevent upload of attachments** to prevent users from transmitting (uploading) attachments to the system.
7. Under Autopolicy: Web Access Control, create a policy that allows or denies users access to the Web resource (and all of its subdirectories) listed in the Resource field.

1. Specify the Web server or HTML page to which you want to control access in the Resource field. Use the format: [protocol://]host[:port] [/path].
2. Select **Allow** to enable access to the specified resource or Deny to block access to the specified resource from the Action list.
3. Click **Add**.

8. Under Autopolicy: Caching, specify the resources to which this policy applies in the Resource box.



The correct caching resource policy must be configured to allow end users to open and save e-mail attachments of different document types in OWA. For example, if the caching policy is set to Smart, end users cannot save .htm or .html attachments to disk.

9. Under Autopolicy: Web Compression, create a policy that specifies which types of Web data the system should and should not compress.

1. Specify the resources to which this policy applies in the Resources box.
2. Select one of the following options from the Action list:
 - **Compress**-Compress the supported content types from the specified resource.
 - **Do not compress**-Do not compress the supported content types from the specified resource.
3. Click **Add**.
10. Select the **Autopolicy: Single Sign-On** check box to pass data such as the username and password to the OWA application.
11. Click **Save and Continue**.
12. Select the roles to which the resource profile applies in the **Roles** tab and click **Add**.
13. The selected roles inherit the autopolicies and bookmarks created by the Microsoft OWA resource profile. If it is not already enabled, the system also automatically enables the Web option in the Users > User Roles > Select _Role > General > Overview page of the admin console.
14. Click **Save Changes**.
15. (Optional.) Modify the default bookmark created by the system in the Bookmarks tab, and/or create new ones.

Microsoft RDWeb HTML5 Templates

Creating Resource Profiles Using the Microsoft RDWeb Template

A Microsoft RDWeb template is a resource profile that controls access to the published desktops and applications based on HTML5. Microsoft RDWeb templates significantly reduce the configuration time by consolidating configuration settings into one place and by pre-populating a variety of resource policy settings.

The following figure depicts Creating Resource Profiles Using the Microsoft RDWeb Template (**Users > Web Application Resource Profiles > New Web Application Resource Profile**):

Web Application Resource Profiles >
New Web Application Resource Profile

Type: *

Name: *

Description:

Base URL: * This URL will be used to create bookmarks to your web application and be used to generate resource policies. We n
 Example: http://www.domain.com

Autopolicy: Web Access Control
 Use this autopolicy to control access to web servers and URLs.

Delete

Resource	Action	
<input type="text"/>	<input type="text" value="Allow"/>	<input type="button" value="Add"/>

Examples:
 http://.domain.com/public/
 https://www.domain.com:443/

To create a resource profile using the Microsoft RDWeb template:

1. Select **Users > Resource Profiles > Web** in the admin console.
2. Click **New Profile**.
3. Select **Microsoft RDWeb** from the **Type** list.
4. Enter a unique name and optionally a description for the Microsoft RDWeb resource profile.

5. Enter the URL of the Microsoft RDWeb resource to which you want to control access in the Base URL field. It is recommended to use the fully qualified domain name with the format: <http://www.domain.com>. The system uses the specified URL to create bookmarks to your web application and be used to generate resource policies.
6. Under **Autopolicy: Web Access Control**, create a policy that allows or denies users access to the web servers and URLs listed in the Resource box:
 1. Specify the Web server or HTML page to which you want to control access in the Resource box. Use the format: [protocol://]host[:port] [/path].
 2. Select **Allow** to enable access to the specified resource or Deny to block access to the specified resource from the Action list.
 3. Click **Add**.
7. Click **Save and Continue**.
8. In the **Roles** tab, select the roles to which the RDWeb resource profile applies and click **Add**. The selected roles inherit the autopolicies and bookmarks created by the RDWeb resource profile.
9. (Optional.) Select the **Bookmarks** tab to modify the default bookmark created by the system and/or create new bookmarks. By default, the system creates a bookmark for the URL defined in the **Base URL** field and displays it to all users assigned to the role specified in the **Roles** tab.

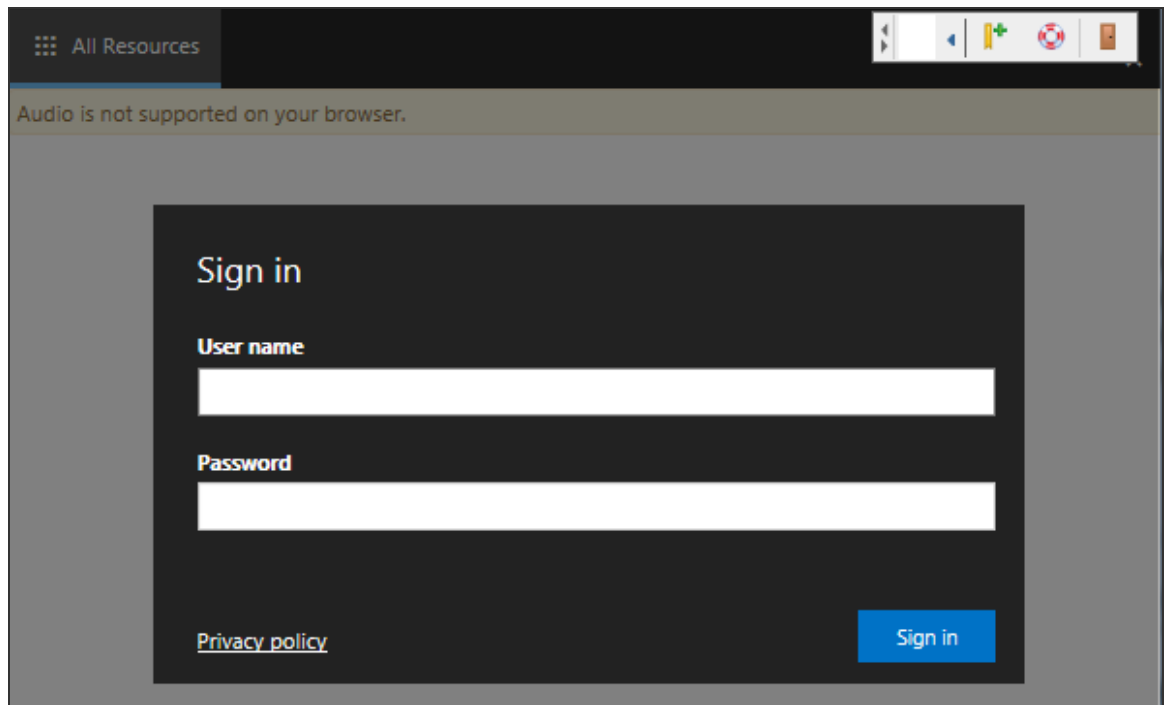
User Experience

With Microsoft RDWeb, a user can launch any application published on RDWeb server using any browser supporting HTML5 technology:

- **Windows 8.1, 10:** Microsoft Edge, Microsoft Internet Explorer, Google Chrome or Mozilla Firefox
- **MacOS:** Safari, Google Chrome, or Mozilla Firefox

1. Log into the Microsoft RDWeb client with username and password.

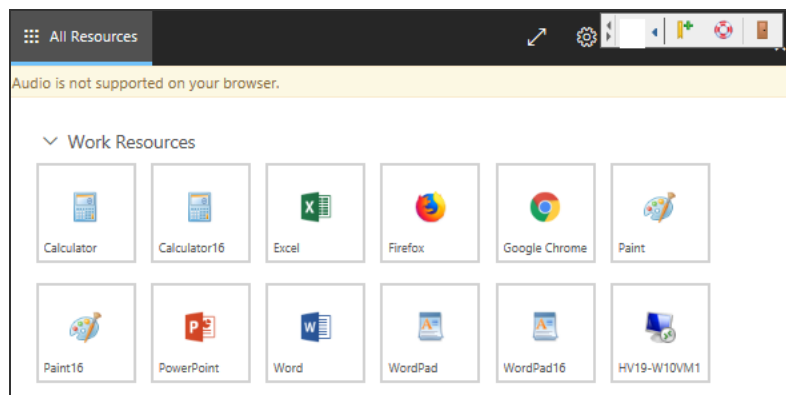
The following figure depicts the Microsoft RDWeb Client Sign in Page:



On successful login, a list of all the applications and virtual desktops published in the RDWeb server is displayed.

2. Click the icon to launch the application.

The following figure depicts the Applications and Virtual Desktops published in the Microsoft RDWeb server:



Microsoft Sharepoint Templates

Creating Resource Profiles Using the Microsoft Sharepoint Template

A Microsoft Sharepoint template is a resource profile that controls access to the application and configures Sharepoint settings as necessary. Microsoft Sharepoint templates significantly reduce your configuration time by consolidating configuration settings into one place and by pre-populating a variety of resource policy settings for you depending on the type of setup you select.

The system supports intermediating traffic to Microsoft Sharepoint through a Web rewriting resource profile template, JSAM, PSAM, and VPN Tunneling. This topic describes how to configure access using the Web rewriting template.

Sharepoint 2010 and Sharepoint 2013 are supported, with the caveats listed in [Table](#).

Sharepoint 2010 and 2013 Caveats.

Version	Caveats
Sharepoint 2010	<p>Office Web Apps through Windows Live is not supported.</p> <p>OneNote document support is limited to only notebooks created to be stored on a local computer and then published on SharePoint 2010.</p> <p>Office documents residing on SharePoint server cannot be opened with Microsoft Office when the server is accessed through the system.</p> <p>In the current release, we support sending contact information from Sharepoint to your Outlook client through the Content Intermediation Engine (Web rewriting feature). Transferring the contact information to the backend Exchange server requires PSAM, JSAM, or VPN Tunneling. To import contact information into the Sharepoint server from your Outlook client, first export your contacts and then upload them to the Sharepoint server.</p>
Sharepoint 2013	<p>Active Directory Federation Services and claims-based authentication are not supported in this release.</p> <p>Integrated service with Exchange 2013 and Lync 2013 is not supported.</p>

To create a resource profile using the Microsoft Sharepoint template:

1. Select **Users > Resource Profiles > Web** in the admin console.
2. Click **New Profile**.
3. Select **Microsoft Sharepoint** from the Type list.
4. Enter a unique name and optionally a description for the Sharepoint resource profile.
5. Enter the URL of the Sharepoint resource to which you want to control access in the Base URL field. Use the format: [protocol://]host[:port][/path]. The system uses the specified URL to define the default bookmark for the Sharepoint resource profile. You may enter a directory URL or a file URL.
6. Under Sharepoint Settings, select **Allow in-line editing of documents within explorer view** to allow users to modify files displayed in the explorer view.



This option is supported only if you enable persistent session (**User > User Roles > RoleName > General > Session Options**.)

7. Enter the URL to the Explorer View page, and then click Add. Do not enter a value that resolves to non-Explorer View URLs (such as http://*:*). Doing so might cause Explorer View to not launch.
8. Order the resources in your list, if appropriate, by selecting the check box next to an item and then using the up and down arrows to move it to the correct place in the list.
9. Enter the number of minutes a persistent cookie resides on a user's computer before it expires in the Persistent cookie timeout box.



Do not confuse this timeout option with Max. Session Length, which determines the number of minutes an active nonadministrative user session may remain open before ending.

10. Select **Add Web ACL** if you have Sharepoint 2013 and the Office Web Apps are on a separate server. The cursor is moved to the Resource text box where you can enter the URL for the Office Web Apps server (see Step 8).
11. Under Autopolicy: Web Access Control, create a policy that allows or denies users access to the Web resource (and all of its subdirectories) listed in the Resource box.

1. Specify the Web server or HTML page to which you want to control access in the Resource box. Use the format: [protocol://]host[:port] [/path].
 2. Select **Allow** to enable access to the specified resource or Deny to block access to the specified resource from the Action list.
 3. Click **Add**.
12. (Optional) Click **Show ALL autopolicy types** to create additional autopolicies that fine-tune access to the resource. Then, create the autopolicies.
13. Click **Save and Continue**.
14. Select the roles to which the resource profile applies in the Roles tab, and click **Add**.
- The selected roles inherit the autopolicies and bookmarks created by the Microsoft Sharepoint resource profile. If it is not already enabled, the system also automatically enables the Web option in the Users > User Roles > Select Role > General > Overview page of the admin console.
15. Click **Save Changes**.
16. (Optional) Modify the default bookmark created by the system in the Bookmarks tab or create new ones.

Web Rewriting

The Web rewriting feature enables you to intermediate Web URLs through the Content Intermediation Engine. You can intermediate URLs on the World Wide Web or on your corporate Intranet. Web rewriting also supports SNI TLS Extension.

When you intermediate standard Web content, you can create supplemental policies that "fine-tune" the access requirements and processing instructions for the intermediated content. You can create these supplemental policies through resource profiles (recommended) or resource policies.

For details about the configuration, refer to the *ICS Web Rewriting Configuration Guide* available on the <https://www.ivanti.com/support/product-documentation> site.

File Rewriting

A file resource profile controls access to resources on Windows server shares or UNIX servers.

When creating a file resource profile, you must use the following formats when defining a resource policy's primary resource as well as its autopolicy resources.

Windows resources: \\server[\share[\path]]

UNIX resources: server[/path]

Within these formats, the three components are:

- **Server (required)** - Possible values:
 - **Hostname** - You may use the system variable <username> when defining the hostname.
 - **IP address** - The IP address needs to be in the format: a.b.c.d

The leading two back slashes are required for Windows, non-Nfs resources.

- **Share (required, Windows only)** - The system variable <username> is allowed. Note that when the system tries to connect to a Windows file share, it connects to ports 445 and 139.
- **Path (optional)** - Special characters allowed include:

*	Matches any character. Note that you cannot use the * wildcard character when defining a resource profile's primary resource (that is, the Server/share field for Windows resources or the Server field for UNIX resources).
%	Matches any character except slash (/)
?	Matches exactly one character

For details about the configuration, refer to the *ICS File Rewriting Configuration Guide* available on the <https://www.ivanti.com/support/product-documentation> site.

PleasSecure Application Manager

Secure Application Manager Overview

The Secure Application Manager option provides secure, application-level remote access to enterprise servers from client applications. You may deploy two versions of the Secure Application Manager:

- Windows version (PSAM) - The Windows version of the Secure Application Manager is a Windows-based solution that enables you to secure traffic to individual client/server applications and application servers.
- Java version (JSAM) - The Java version of the Secure Application Manager provides support for static TCP port client/server applications, including enhanced support for Microsoft MAPI, Lotus Notes, and Citrix NFuse. JSAM also provides NetBIOS support, which enables users to map drives to specified protected resources.



Beginning ICS 9.0R4, support is not available for the legacy Pulse Windows Secure Application Manager (PSAM) clients. Users are recommended to migrate to Pulse Desktop Client to get continued maintenance. For migration details, refer to the *PSAM to Pulse SAM Migration Guide*.


Task Summary: Configuring PSAM

This section provides high-level PSAM configuration steps. These steps do not account for preliminary system configuration steps such as specifying the system's network identity or adding user IDs.

To configure PSAM:

1. Create resource profiles that enable access to client/server applications or destination networks, create supporting autopolicies as necessary, and assign the policies to user roles using settings in the Users > Resource Profiles > SAM pages of the admin console.

The following figure depicts Configuring PSAM (**Users > User Roles > <test1> > General > Overview**):



We recommend that you use resource profiles to configure PSAM (as described above). However, if you do not want to use resource profiles, you can configure PSAM using role and resource policy settings in the following pages of the admin console instead:

- Enable access to PSAM at the role-level using settings in the **Users > User Roles > Role > General > Overview** page of the admin console.

- Specify which client/server applications and servers PSAM should intermediate using settings in the **Users > User Roles > SAM > Applications** page of the admin console.
 - Specify which application servers' users can access through PSAM using settings in the **Users > Resource Policies > SAM > Access** page of the admin console.
2. After enabling access to client/server applications and/or destination networks using PSAM resource profiles or roles and resource policies, you can modify general role and resource options in the following pages of the admin console:
 - (Optional) Configure role-level options such as whether the system should automatically launch and upgrade PSAM using settings in the **Users > User Roles > SAM > Options** page of the admin console.
 - (Optional) Control IP based hostname matching at the resource level using settings in the **Users > Resource Policies > SAM > Options** page of the admin console.
 3. Ensure that an appropriate version of PSAM is available to remote clients using settings in the **Maintenance > System > Installers** page of the admin console.
 4. If you want to enable or disable client-side logging for PSAM, configure the appropriate options through the **System > Configuration > Security > Client-side Logs** tab of the admin console.

PSAM Recommended Operation

Ivanti recommends the following operation when using PSAM:

- PSAM supports client-initiated UDP and TCP traffic by process name, by destination hostname, or by destination address range:port range. Except for Passive FTP, PSAM only supports protocols that do not embed IP addresses in the header or payload. W-SAM also supports unicast client-initiated UDP.
- Users must launch drive maps through PSAM in one of the following ways:
 - NetUse - At the Command prompt, type **net use * \\server\share /user:username.**
 - Right-click My Computer > Map Network Drive, or in Windows Explorer, go to **Tools > Map Network Drive and select Connect using a different username.**

- When using the PSAM Access Control List (ACL), administrators should take extra precaution when granting access to hosts. We recommend that administrators use the IP address instead of the hostname. If the hostname is required for security purposes, administrators should try to include additional ACLs with the corresponding IP address or IP addresses for that hostname. Reverse DNS lookups are not supported.
- To run Citrix NFuse through W-SAM, you must define a Caching rule to cache launch.asp files. For example, configure the resource policy to *server name:80,443/*launch.asp* and the **Caching Option to Cache (do not add/modify caching headers)**.
- When using PSAM on Pocket PC, session roaming should be enabled when being used over GPRS because the IP address of the phone may change.
- When using PSAM on Pocket PC, if you have multiple roles defined, select the **Merge settings for all assigned roles** option under Administrators > Admin Realms > Realm > Role Mapping.
- When using an external load balancer and accessing J-SAM, W-SAM, persistence must be employed on the load balancer. This persistence should be based on Source IP or Destination Source, depending on the load balancer being used.

Debugging PSAM Issues

You can use the Secure Application Manager dialog box on an end-user's system to view the PSAM status and a variety of details about the user's session. For instance, the Secure Application Manager dialog box displays the applications and servers that PSAM is configured to secure, event logs and Winsock data for the user's session, and various system diagnostics and performance data. This information can help you or a Ivanti Support representative debug any problems your users may encounter.

To access the Secure Application Manager dialog box, users simply need to double-click the PSAM icon on their Windows task bars:



For more information about viewing information in the Secure Application Manager dialog box, see the end-user help system available from the Help link in the end-user console.

About PSAM Resource Profiles

You can create two types of PSAM resource profiles:

- PSAM application resource profiles-These resource profiles configure PSAM to secure traffic to a client/server application. When you create a PSAM application resource profile, the PSAM client intercepts requests from the specified client applications to servers in your internal network.
- PSAM destination network resource profiles-These resource profiles configure PSAM to secure traffic to a server. When you create a PSAM destination network resource profile, the PSAM client intercepts requests from processes running on the client that are connecting to the specified internal hosts.

When creating PSAM resource profiles, note that the resource profiles do not contain bookmarks. To access the applications and servers that PSAM intermediates, users must first launch PSAM and then launch the specified application or server using standard methods (such as the Windows Start menu or a desktop icon).

When you enable JSAM or PSAM through Web rewriting autopolicies in the Users > Resource Profiles > Web Applications/Pages page of the admin console, the system automatically creates JSAM or PSAM autopolicies for you. You can only view these SAM policies through the appropriate Web resource profile-not through the SAM resource profile pages of the admin console.

Creating PSAM Client Application Resource Profiles

When you create a PSAM application resource profile, the PSAM client intercepts requests from the specified client applications to servers in your internal network.

To create a PSAM application resource profile:

1. In the admin console, choose **Users > Resource Profiles > SAM > Client Applications**.

The following figure depicts Creating PSAM Client Application Resource Profiles (**Users > Client Application Resource Profiles > New Client Application Resource Profile**):

Client Application Resource Profiles >
New Client Application Resource Profile

Type: * WSAM
 Application: * Domain Authentication
 Name: *
 Description:

Add domain controller(s) into the WSAM Destinations. Kerberos uses port 88 and LDAP uses port 389 by default. DNS SRV requests will be auto-allowed from WSAM.

Autopolicy: SAM Access Control
 Use this policy to control access to application servers.
 Resources:

Resource	Action	
<input type="text"/>	Allow	<input type="button" value="Add"/>

Resource Examples: <USER>.domain.dom:22,23
 appserver*.domain.com:*
 10.10.10.10/255.255.0:80,443,8080
 10.10.10.10/24:8000-9000

*indicates required field

2. Click **New Profile**.
3. From the Type list, choose **PSAM**.
4. From the Application list, select one of the following options:
 - **Custom**-When you select this option, you must manually enter your custom application's executable file name (such as telnet.exe). Additionally, you may specify this file's path and MD5 hash of the executable file (although it is not required that you specify the exact path to the executable). If you enter an MD5 hash value, PSAM verifies that the checksum value of the executable matches this value. If the values do not match, PSAM notifies the user that the identity of the application could not be verified and does not forward connections from the application to the system.
 - **Lotus Notes**-When you select this option, PSAM intermediates traffic from the Lotus Notes fat client application.
 - **Microsoft Outlook**-When you select this option, PSAM intermediates traffic from the Microsoft Outlook application.
 - **NetBIOS file browsing**-When you select this option, PSAM intercepts NetBIOS name lookups in the TDI drivers on port 137.
 - **Citrix**-When you select this option, PSAM intermediates traffic from Citrix applications.

You can only use PSAM to configure access to a standard application once per user role. For example, you can enable one configuration of Microsoft Outlook and one configuration of Lotus Notes for the "Users" role.

The system supports several mechanisms for intermediating traffic to the Lotus Notes, Microsoft Outlook, and Citrix applications.

- Domain Authentication-Select this option to allow integrated Windows applications, such as file sharing, Outlook, and so forth to authenticate to the domain controller when the client machine is part of a domain. Before using this option, you must:
 - Specify domain controllers that are reachable through the system in the PSAM Destination list so that LDAP and Kerberos traffic can be proxied and sent to the system.
 - Configure a PSAM Access Control Policy (ACL) to allow access to all domain controllers.
5. Enter a unique name and optionally a description for the resource profile. The system displays this information in the Client Application Sessions section of the end-user home page.
 6. In the **Autopolicy: SAM Access Control** section, create a policy that allows or denies users access to the server that hosts the specified application.

The following figure depicts Autopolicy: SAM Access Control:

Autopolicy: SAM Access Control

Use this policy to control access to application servers.

Resources: Delete ↑ ↓

<input type="checkbox"/>	Resource	Action	
	<input type="text"/>	Allow	Add

Resource Examples: <USER>.domain.dom:22,23
 appserver*.domain.com:*
 10.10.10.10/255.255.255.0:80,443,8080
 10.10.10.10/24:8000-9000

Save and Continue >

*indicates required field

7. If it is not already enabled, select the **Autopolicy: SAM Access Control** check box.
8. In the Resource field, specify the application server to which this policy applies. You can specify the server as a hostname or an IP/netmask pair. You may also include a port.

If you select Domain Authentication from the Application list, enter your domain controller server addresses into the Resource field. You can add multiple domain controller servers if more than one is available.

When enabling auto-policy for any client application for PSAM, avoid entering *.* in the resource list since the access control policies are not restricted to that particular application. This may result in other resources being accessed through client applications for which the access control policies are not defined.

9. From the Action list, select **Allow** to enable access to the specified server or **Deny** to block access to the specified server.
10. Click **Add**.
11. Click **Save** and **Continue**.
12. In the Roles tab, select the roles to which the resource profile applies and click **Add**.

The selected roles inherit the autopolicy created by the resource profile. If it is not already enabled, the system also automatically enables the SAM option in the Users > User Roles > Select Role > General > Overview page of the admin console for all of the roles you select.

13. Click **Save Changes**.

Creating PSAM Destination Network Resource Profiles

When you create a PSAM destination network resource profile, the PSAM client intercepts requests from processes running on the client to internal hosts.

When destinations (using either IP address or hostnames) are configured on the system, all DNS and NetBIOS names are resolved through the system.

To create a PSAM destination network resource profile:

1. In the admin console, choose **Users > Resource Profiles > SAM > PSAM Destinations**.
2. Click **New Profile**.
3. Enter a unique name and optionally a description for the resource profile.
4. In the PSAM Destinations section, specify which servers you want to secure using PSAM and click **Add**. You can specify the servers as hostname or IP/netmask pairs. You may also include a port.

5. Select the Create an access control policy allowing SAM access to this server check box to enable access to the server specified in the previous step (enabled by default).
6. Click **Save** and **Continue**.
7. In the Roles tab, select the roles to which the resource profile applies and click Add.

The selected roles inherit the autopolicy created by the resource profile. If it is not already enabled, the system also automatically enables the SAM option in the Users > User Roles > Role Name > General > Overview page of the admin console for all of the roles you select.

Specifying Applications and Servers for PSAM to Secure

Information in this section is provided for backwards compatibility. We recommend that you secure traffic using PSAM resource profiles instead, since they provide a simpler, more unified configuration method.

Use the Applications tab to specify applications and servers for which PSAM secures traffic. When PSAM downloads to a client PC, it contains the information you configure on the Applications tab for the role. After a user launches the Secure Application Manager, PSAM intercepts requests from client applications to servers in your internal network and requests from processes running on the client to internal hosts. You define these resources on the Applications tab by configuring two lists:

- PSAM supported applications list-This list contains applications for which you want PSAM to secure client/server traffic between the client and the system.
- PSAM allowed servers list-This list contains hosts for which you want PSAM to secure client/server traffic between the client and the system.

To specify applications for which PSAM secures client/server traffic between the client and the system:

1. In the admin console, choose **Users > User Roles > Select Role > SAM > Applications**.
2. Click **Add Application**.
3. Enter the name of the application and, optionally, a description. This information displays in the Client Application Sessions section of the end-user home page.
4. From the Type list, choose one of the following options:
 - **Standard**-If you select this option, choose one the following applications from the Application Parameters section:
 - **Citrix**-When you select this option, PSAM intermediates traffic from Citrix applications.

- **Lotus Notes**-When you select this option, PSAM intermediates traffic from the Lotus Notes fat client application.
- **Microsoft Outlook/Exchange**-When you select this option, PSAM intermediates traffic from the Microsoft Outlook application.
- **NetBIOS file browsing**-When you select this option, PSAM intercepts NetBIOS name lookups in the TDI drivers on port 137.

Note that in order to access a share using PSAM with NetBIOS, you need to explicitly specify the server's NetBIOS name (alphanumeric string up to 15 characters) in two places: on the Add Server page and in a SAM resource policy. (Wildcards are currently not supported.) Alternatively, you can enable the Auto-allow application servers option on the SAM > Options tab, and then the system automatically creates a SAM resource policy that allows access to this server.

- **Custom**-Select this option to specify a custom client/server application. Then:
 - In the Filename field, specify the name of the file's executable file.
 - Optionally specify the file's path and MD5 hash of the executable file. If you enter an MD5 hash value, PSAM verifies that the checksum value of the executable matches this value. If the values do not match, PSAM notifies the user that the identity of the application could not be verified and does not forward connections from the application to the system.
5. Click **Save Changes** or **Save + New**.
 6. Configure a PSAM resource policy to specify to which enterprise resources (based on IP address/port combination) the system may send the application.

Specifying Servers for PSAM to Secure

To specify servers for which PSAM secures client/server traffic between the client and the system:

1. In the admin console, choose **Users > User Roles > Select Role > SAM > Applications**.
2. Click **Add Server**.
3. Enter the name of the server and, optionally, a description.
4. Specify the server's hostname (the wild cards '*' or '?' are accepted) or an IP/netmask pair. Specify multiple ports for a host as separate entries.
5. Click **Save Changes** or **Save + New**.

6. Configure a PSAM resource policy to specify to which enterprise resources (based on IP address/port combination) the system may send a server request.

Alternatively, you can enable the Auto-allow application servers option on the SAM > Options tab, and then the system automatically creates a SAM resource policy that allows access to the specified server. Note that you need to enable this option before specifying the application or server; otherwise, you need to create a SAM resource policy.

Specifying Applications that Need to Bypass PSAM

The PSAM client comes pre-configured with a list of "passthrough" applications bypass PSA. The PSAM client does not secure traffic for these applications. In addition to bypassing these predefined applications, you may also specify additional applications that should bypass PSAM.



PSAM does not bypass applications on Pocket PCs and other handheld devices.

To specify applications for PSAM to secure:

1. In the admin console, choose **Users > User Roles > Select Role > SAM > Applications**.
2. Click **Add Bypass Application**. The New Bypass Application page displays.
3. Name the application and provide a description (optional).
4. Provide the file name (required).
5. Enter the absolute path to the application (optional).
6. Select **Save Changes** to add the bypass application to the list or **Save + New** to save the bypass application and create another bypass application.

Default Bypass Applications

The PSAM client is preconfigured to bypass PSAM processing for the following applications:

- apache.exe
- apache*
- licadmin.exe
- vni.exe
- lmgrd.exe
- TNSLSNR.EXE

- ORACLE.EXE
- Agntsvr.exe
- ONRSD.EXE
- Pagntsvr.exe
- ENCSVC.EXE
- Agntsvc.exe
- EiSQLW.exe
- Sqlservr.exe
- Sqlmangr.exe
- inetinfo.EXE
- xstart.exe
- idsd.exe
- dsTermServ.exe
- dsCitrixProxy.exe
- dsNcService.exe
- dsNetworkConnect.exe

Specifying Role-Level PSAM Options

To specify PSAM options at the role level:

1. In the admin console, choose **Users > User Roles > Select Role > SAM > Options**.
2. If it is not already enabled, select the Windows SAM option at the top of the page.
3. Under Secure Application Manager options, configure the following options:
 - **Auto-launch Secure Application Manager**-If you enable this option, the system automatically launches the Secure Application Manager when a user signs in. If you do not select this option, users must manually start the Secure Application Manager from the Client Applications Sessions section of the end-user home page.

Although you configure the Secure Application Manager to automatically launch when users sign into the device, users can override this setting through the Preferences > Applications page of the end-user console. If you or the end user disables PSAM from automatically launching, users need to manually start the Secure Application Manager by clicking its link on the home page.

- **Auto-allow application servers**-If you enable this option, the system automatically creates a SAM resource policy that allows access to the server specified in the PSAM application and server lists.

You may not see the Auto-allow option if you are using a new installation or if an administrator hides the option.

4. Under Windows SAM Options, configure the following options:

- **Auto-uninstall Secure Application Manager**-Select this option to automatically uninstall the Secure Application Manager after users sign off.
- **Prompt for username and password for intranet sites**-Select this option to require users to enter their sign-in credentials before connecting to sites on your internal network. This option changes Internet Explorer's intranet zone setting so that Internet Explorer prompts the user for network sign-in credentials whenever the user wants to access an intranet site.
- **Auto-upgrade Secure Application Manager**-Select this option to automatically download the Secure Application Manager to a client machine when the version of Secure Application Manager on the system is newer than the version installed on the client. If you select this option, note the following:
 - The user must have Administrator privileges in order for the system to automatically install Secure Application Manager on the client.
 - If a user uninstalls Secure Application Manager and then signs in to the system for which the Auto-upgrade Secure Application Manager option is not enabled, the user no longer has access to Secure Application Manager.
- **Resolve only hostnames with domain suffixes in the device DNS domains**-If this option is configured, PSAM filters DNS requests (FQDNs) and sends to the system only those DNS requests that have a domain suffix in the list of DNS Domains configured on the Network Overview page. This option is limited to resolution of FQDNs only. No filtering is applied to short names and NetBIOS requests.

- **Session start script and Session end script**-If you want to run a batch, application, or Win32 service file when the PSAM session starts or ends, enter the name and path for the file. For example, if you want to terminate an application and then restart it, you may use PSKILL.exe (a third-party utility that terminates processes on local or remote systems).

If you enable the Session start script option or Session end script option, note the following:

- You must either install the specified file on your end-user's computers or specify a path on an accessible network directory.
 - To ensure that the system can locate a file on different platforms, you can use Windows variables, such as in a path such as %WINDIR%\system32\log.
 - The file must invoke the PSAM launcher using the appropriate command-line options.
5. Click Save Changes.

Specifying Application Servers that Users Can Access

Information in this section is provided for backwards compatibility. We recommend that you secure traffic using PSAM resource profiles instead, since they provide a simpler, more unified configuration method.

When you enable the Secure Application Manager access feature for a role, you need to create resource policies that specify which application servers a user may access. These policies apply to both the Java version and Windows version of the Secure Application Manager (JSAM and PSAM, respectively). When a user makes a request to an application server, the system evaluates the SAM resource policies. If the system matches a user's request to a resource listed in a SAM policy, the system performs the action specified for the resource.

When writing a SAM resource policy, you need to supply key information:

- **Resources**-A resource policy must specify one or more resources to which the policy applies. When writing a SAM policy, you need to specify application servers to which a user may connect.
- **Roles**-A resource policy must specify the roles to which it applies. When a user makes a request, the system determines what policies apply to the role and then evaluates those policies that correspond to the request. SAM resource policies apply to users' requests made through either version, JSAM or PSAM.
- **Actions**-A Secure Application Manager resource policy either allows or denies access to an application server.

You can create resource policies through the standard interface (as described in this section) or through resource profiles (recommended method).

The engine that evaluates resource policies requires that the resources listed in a policy's Resources list follow a canonical format.

To write a Secure Application Manager resource policy:

1. In the admin console, choose **Users > Resource Policies > SAM > Access**.
2. On the Secure Application Manager Policies page, click **New Policy**.
3. Enter a name to label this policy (required) and a description of the policy (optional).
4. In the Resources section, specify the application servers to which this policy applies.
5. In the Roles section, specify:
 - **Policy applies to ALL roles**-Choose this option to apply this policy to all users.
 - **Policy applies to SELECTED roles**-Choose this option to apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
 - **Policy applies to all roles OTHER THAN those selected below**-Choose this option to apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
6. In the Action section, specify:
 - **Allow socket access**-Choose this option to grant access to the application servers specified in the Resources list.
 - **Deny socket access**-Choose this option to deny access to the application servers specified in the Resources list.
 - **Use Detailed Rules**-Choose this option to specify one or more detailed rules for this policy.
7. Click **Save Changes**.
8. On the Secure Application Manager Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Specifying Resource Level PSAM Options

Use the Options tab to specify the SAM resource option to match IP addresses to hostnames specified as resources in your SAM resource policies. When you enable this option, the system looks up IP addresses corresponding to each hostname specified in a SAM resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

When you enable this option, the system compiles a list of hostnames specified in the Resources field of each SAM resource policy. The system then applies the option to this comprehensive list of hostnames.



This option does not apply to hostnames that include wildcards and parameters.

To specify the SAM resource option:

1. In the admin console, choose **Users > Resource Policies > SAM > Options**.
2. Select **IP based matching for Hostname based policy resources**. This option looks up the IP address corresponding to each hostname specified in a Secure Application Manager resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.
3. Click **Save Changes**.

JSAM Overview

The Java version of the Secure Application Manager provides support for static TCP port client/server applications, including enhanced support for Microsoft MAPI, Lotus Notes, and Citrix NFuse. JSAM also provides NetBIOS support, which enables users to map drives to specified protected resources.

JSAM works well in many network configurations but does not support dynamic port TCP-based client/server applications, server-initiated connections, or UDP traffic.



regedit.exe is required for some JSAM functionality. If regedit.exe is disabled, automatic host mapping and the NetBIOS and Outlook/Exchange applications will not work properly.

For information about the operating systems, Web browsers, and JVMs on which Ivanti supports JSAM, see the *Supported Platforms Guide*.

Task Summary: Configuring JSAM

This topic provides high-level JSAM configuration steps. These steps do not account for preliminary system configuration steps such as specifying the system's network identity or adding user IDs.

To configure JSAM:

1. Create resource profiles that enable access to client/server applications, create supporting autopolicies as necessary, and assign the policies to user roles using settings in the Users > Resource Profiles > SAM pages of the admin console.

We recommend that you use resource profiles to configure JSAM (as described above). However, if you do not want to use resource profiles, you can configure JSAM using role and resource policy settings in the following pages of the admin console instead:

1. Enable access to JSAM at the role-level using settings in the Users > User Roles > Select Role > General > Overview page of the admin console.
 2. Specify which client/server applications JSAM should intermediate using settings in the Users > User Roles > SAM > Applications page of the admin console.
 3. Specify which application servers' users can access through JSAM using settings in the Users > Resource Policies > SAM > Access page of the admin console.
2. After enabling access to client/server applications using JSAM resource profiles or roles and resource policies, you can modify general role and resource options in the following pages of the admin console:
 - (Optional) Configure role-level options such as whether the system should automatically launch JSAM using settings in the Users > User Roles > SAM > Options page of the admin console.
 - (Optional) Control IP based hostname matching at the resource level using settings in the Users > Resource Policies > SAM > Access page of the admin console.
 3. If you want to enable or disable client-side logging for JSAM, configure the appropriate options through the System > Configuration > Security > Client-side Logs tab of the admin console.

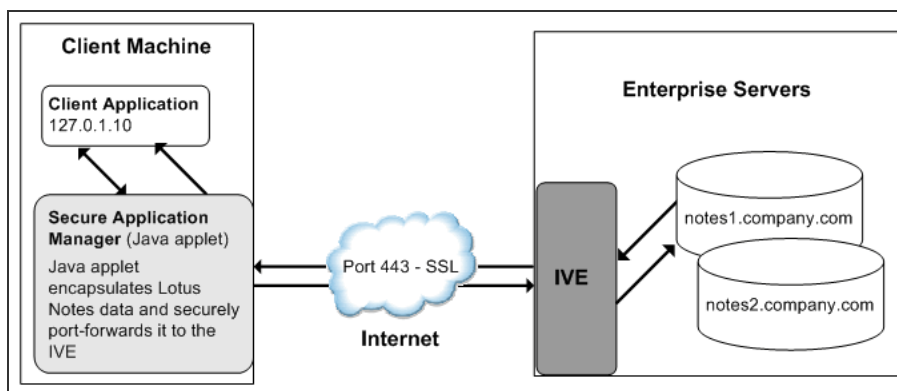
4. If you have multiple internal domains, such as company-a.com and company-b.com, add DNS domains to the system using settings in the System > Network > Overview page of the admin console so that names such as app1.company-a.com and app2.company-b.com resolve correctly.
5. If a remote user's PC is set up to use a Web proxy in Internet Explorer, configure the client machine to bypass the proxy server when the user launches applications that need to connect to the Secure Application Manager.
6. Enable JSAM to associate IP loopback addresses with application servers on specific ports either by enabling JSAM to edit the hosts file on your users' systems or by creating an external DNS to route client application traffic to the JSAM applet.

Using JSAM for Client/Server Communications

JSAM provides secure port forwarding by directing client application traffic to the JSAM applet running on a client machine. To the client application running on the local machine, JSAM appears as the application server. To the application server in your network, the system appears as the client application.

The below figure illustrates the interaction between a client application and its server via Ivanti Connect Secure. (This figure assumes that the user specified a localhost IP address as the server in the client application.)

The following figure depicts the Java Secure Application Manager:




1. The user starts a client application listed in the Client Application Sessions section of the end-user home page. The application resolves the remote server to localhost.
2. The client application connects to JSAM running on the user's machine and starts sending requests.
3. JSAM encapsulates and forwards all client requests to the system over SSL.

4. The system unencapsulates the client data and forwards it to the specified application server.
5. The application server responds with data to the system.
6. The system encapsulates and forwards the response from the application server to JSAM over SSL.
7. JSAM unencapsulates the application server data and forwards it to the client application.

A status indicator on the JSAM window shows the current state of JSAM. If green, JSAM is working correctly. If red, JSAM is unable to send/receive requests to/from the system.

The JSAM window updates the status indicator only when traffic is passed through JSAM. If no traffic is passed through JSAM, the status indicator remains in its current state. For example, if there is a network outage or if the user's session times out, the status indicator remains green even though it cannot send/receive requests to/from the system.

- If a remote user's PC is set up to use a Web proxy in Internet Explorer, you must configure the client machine to bypass the proxy server when the user launches applications that need to connect to the Secure Application Manager.
 - JSAM allocates 20-30 MB of RAM when running (the exact amount of memory depends on the Java Virtual Machine (JVM) used) and, if caching is enabled, may leave a .jar file on the client machine. For more information about files left by JSAM on client machines, see the Client-side Changes Guide on the Global Support Center.
 -  Users may experience problems waiting for the Secure Application Manager to fully load if they enable pop-up blockers through their Web browsers. This problem occurs because a pop-up window alerting users to accept the Secure Application Manager plug-in may appear in the background (behind the Web browser window) where users cannot see it.
 - When launching applications through JSAM, Ivanti supports configuration of 1200 unique IP/port combinations on Windows and Mac and 800 unique IP/port combinations on Linux. Note that this limit is based on IP/port combinations, not applications (which may listen on more than one IP address and port). Ivanti determined these numbers by testing on Windows machines using default JRE memory settings.
-

Assigning IP Loopback Addresses to Servers

For JSAM to function, it must listen on loopback addresses for client requests to network application servers. The system assigns these unique IP loopback address to each application server that you specify for a given port. For example, if you specify:

app1.mycompany.com, app2.mycompany.com. app3.mycompany.com,...

for a single port, the system assigns a unique IP loopback address to each application:

127.0.1.10, 127.0.1.11, 127.0.1.12,...

When the system installs JSAM on a user's machine, JSAM listens on the loopback addresses (on the corresponding client port specified for the application server) for client requests to network application servers. You can configure the system to dynamically assign these loopback addresses, or you can configure static loopback addresses yourself through the admin console.

You must enable these associations between IP loopback addresses and applications servers on a specific port in one of two ways:

- Allow the system to edit the hosts file on the client system with IP loopback assignments. The system makes a copy of the current hosts file and then creates a new hosts file with the IP loopback assignments. When the user ends the session, the system deletes the new hosts file and restores the original hosts file.

If the client system shuts down unexpectedly, the hosts file still points the client to loopback addresses for outside connections. Settings in the hosts file are returned to their original state when the client system reboots.

Users must have the proper privileges on their machines in order for the system to edit the hosts file.

- Create an external DNS to route client application traffic to the JSAM applet.

Using Static Loopback Addresses

Using an external DNS server with dynamic loopback addresses requires an administrator to update the DNS settings each time the JSAM application configuration changes. On the other hand, configuring an external DNS server using static loopback addresses provides administrators with the highest degree of configuration control. For example, consider the following IP loopback assignments:

```
app1.mycompany.com - 127.0.1.10
app2.mycompany.com - 127.0.1.11
app3.mycompany.com - 127.0.1.12
```

If you configure an external DNS server using dynamic loopback address assignments and you delete the first application server, the address assignments change:

```
app2.mycompany.com - 127.0.1.10
app3.mycompany.com - 127.0.1.11
```

With static IP loopback addresses in an external DNS, deleting the first application server does not affect the IP loopback assignments for the remaining application servers:

```
app2.mycompany.com - 127.0.1.11
app3.mycompany.com - 127.0.1.12
```

You can assign static IP loopback addresses when creating a JSAM custom resource profile through the Users > Resource Profiles > SAM > Client Applications page of the admin console or when enabling JSAM applications through the Users > User Roles > Select Role > SAM > Applications page of the admin console.

If you assign a static IP loopback address while creating a new application, the system checks the address for conflicts against other configured applications in the same role. If another application uses the same address, the system displays an error message prompting you to enter a different IP address.



Static IP loopback addresses apply only to application servers configured by an administrator. The system assigns dynamic IP loopback addresses for user-defined application servers. If the administrator does not assign an IP loopback address to an application server, the system assigns a dynamic address.

IP Loopback Address Considerations When Merging Roles

- IP Loopback Address Considerations When Merging Roles
- If two or more roles map to the same application and each mapping contains a different static IP loopback address, all of the static IP loopback addresses remain unchanged.
- If two or more roles map to the same application and only one role uses a static IP loopback address, JSAM uses only the static IP loopback address and binds to only one statically defined socket on the client.
- If two or more roles map to the same application using dynamic IP loopback addresses, only one dynamic IP loopback address is used. The application listener binds to only one dynamically assigned socket on the client.
- If you use the same hostname in multiple roles, either use the same static IP loopback address, or dynamic addresses for all the applications.

- If you use different hostnames associated with the same loopback address and port combination, JSAM cannot distinguish between the two different hosts at the back-end and, hence, cannot accurately direct IP traffic bound for those hosts.

Resolving Hostnames to Localhost

For JSAM to successfully intermediate traffic, a client application on the user's machine needs to resolve the application server to the client localhost. This process enables JSAM to capture and securely port forward the data intended for the application server via Ivanti Connect Secure. JSAM can perform automatic host-mapping, in which it edits the client's hosts file, to map application servers to localhost. (You can enable automatic host-mapping through the Users > User Roles > Select Role > SAM > Options page of the admin console.)

In order for JSAM to edit a user's hosts file, the user must have the appropriate authority on the client machine:

- Windows users using the FAT file system may belong to any user group. For Exchange MAPI support, however, users must have at least Power User privileges on their machines.
- Windows users using the NTFS file system must have Administrator privileges on their machines.
- Linux (RedHat) users must launch the browser that will launch JSAM as root.
- Macintosh users must supply the Administrator password when prompted by JSAM.
- If users do not have the appropriate privileges on their machines, JSAM cannot automatically edit the hosts file, preventing hostname resolution to localhost.

Alternatives for users who do not have the appropriate privileges are:

- You configure your external DNS server to resolve application servers to localhost. If you configure your external DNS server to use a localhost address instead of the application server hostname, remote users need to configure the order in which their machine searches DNS servers to start with the corporate DNS.
- You relax the permissions on the etc directory and the etc\hosts file to enable JSAM to make the necessary modifications.
- Users configure a client application to use the localhost address assigned by the system where they typically specify the application server hostname in the client application.

Configuring a PC that Connects Through a Proxy Web Server

If a remote user's PC is set up to use a Web proxy in Internet Explorer, you must configure the client machine to bypass the proxy server and contact the Secure Application Manager instead.

To configure a PC that connects to the system through a Web proxy in Internet Explorer:

1. From the Internet Explorer Tools menu, choose **Internet Options**.
2. On the Connections tab, click the **LAN Settings** button.
3. Under Proxy server, click the **Advanced** button.
4. Under Exceptions, enter the addresses for which you do not want to use a proxy server. Enter all addresses (hostnames and localhost) that the client application uses when connecting through the Secure Application Manager. For example:

If your application server is app1.company.com, enter the following exceptions:

```
app1;app1.company.com;127.0.0.1
```

If your Exchange Server is exchange.company.com, enter the following exceptions:

```
exchange;exchange.company.com;127.0.0.1
```



Connect Secure clients parse Internet Explorer's static proxy exception list. We support most exceptions that Internet Explorer supports with the following limitations:

- For IP address exception, we support n.*.*, n.n.*, n.n.n*. For example, 10.*.*, 10.10.*, 10.10.10.*, or 10.10.10.10. We do not support 10* or 10*.10.* even though Internet Explorer may support them.
- For string expression, we support specific strings such as my.company.net, or a wild card at front of the string, for example, *.my.company.net or *.company.net. We do not support *.company.*, *.company*, *.company*.com, *.net, *.com and so forth.

Determining the Assigned Loopback Address

Users cannot modify the corporate DNS server for applications they add for port forwarding. If you allow users to specify applications for JSAM to proxy, users need to configure a client application to use the localhost address assigned by the system where they typically enter the server hostname.

The Details pane of the JSAM browser window displays the loopback IP address assigned by the system along with the port specified by the user. To determine what IP address the system assigns to an application specified through the Client Applications page, a user must restart the Secure Application Manager after adding the application. The loopback address assigned to the application appears on the Details pane of the Secure Application Manager browser window.

In the client application, the user needs to enter the system-assigned loopback address as the application server. For example, if a user wants to access a telnet server behind your corporate firewall, the user needs to follow these steps:

1. In the Client Application Sessions section of the end-user home page, click the Item Properties icon, then click **Add Application**
2. On the Add Application page, specify:
 - The server's fully qualified domain name or IP address in the Remote Server field, such as terminalserver.pulsesecure.net.
 - The port on which JSAM should listen for client traffic to the server in the Client Port field, such as 3389.
 - The port on which the remote server should listen for traffic from the client application (JSAM) in the Server Port field, such as 3389.
3. Click **Add** to save the information.
4. Close the Secure Application Manager browser window.
5. In the Client Application Sessions section of the end-user home page, click Start to restart the Secure Application Manager.
6. In the Secure Application Manager browser window, click Details.
7. On the Details tab, look at which loopback address is assigned to the remote server, such as 127.0.1.18.
8. In the client application, such as Remote Desktop Connection, specify the loopback address in the configuration field for the server. This field appears in different places for different applications. Users may enter this information through a setup wizard or other configuration dialog.

Configuring External DNS Servers and User Machines

Client applications must resolve server hostnames to JSAM, which proxies data between a client and a server. On Windows PCs, server hostnames are stored in the hosts file. To intercept data using JSAM, the server names in the hosts file need to resolve to the local machine (localhost) so that the system can intermediate the traffic. The recommended process for mapping application servers to a user's local PC is to enable the automatic host-mapping option, which enables the system to automatically modify the PC hosts file to point application servers to the localhost for secure port forwarding.

For the system to perform automatic host-mapping, however, PC users must have the proper privileges on their machines. If your PC users do not have these privileges, you must ensure that your internal application server names resolve externally to a PC's localhost by adding entries to your external Internet-facing DNS server such as:

```
127.0.0.1 app1.company-a.com
127.0.0.1 app2.company-b.com
127.0.0.1 exchange1.company-a.com
127.0.0.1 exchange1.company-b.com
```

If the client application uses an unqualified name for the application server, users need to specify DNS suffixes so that the PC can attach the suffix and contact your external DNS server for name resolution. For example, an MS Outlook client typically has an unqualified name for an MS Exchange server. In order for the qualified name to resolve to 127.0.0.1, users need to specify the appropriate DNS suffixes on their PCs. Adding domain names does not affect other operations on the PC, including use of the client application from within the enterprise.

To configure a user PC with DNS suffixes:

1. From the Windows Start menu, choose **Settings > Network and Dial-up Connections > Local Area Connection** and then choose **Properties**.
2. Select **Internet Protocol (TCP/IP)** and then click **Properties**.
3. Click **Advanced** and then click the **DNS** tab.
4. Click **Append these DNS suffixes** and then click **Add**.
5. Add your enterprise's internal domains as additional DNS suffixes.

JSAM Linux and Macintosh Support

Linux users do not have access to ports below 1024 unless they are signed into their machines as root. Macintosh users do not have access to ports below 1024 unless they supply the Administrator password when prompted by JSAM. To support applications that run on privileged ports (ports below 1024), such as a telnet application:

- Users may launch the browser that will launch JSAM as root.
- You or the user may specify a client port number equal to or greater than port 1024 when enabling client applications.

For example, if you specify 2041 for the client port and 23 for the server port for a telnet application, the command to run the application is:

```
telnet loopbackIP 2041
```

where loopbackIP is the loopback IP address assigned to the application server by the system. JSAM listens on port 2041 for traffic from the telnet application and forwards it to the system. The system then forwards the traffic to port 23 on the destination server.



Due to the design of the Sun JVM code, Macintosh users cannot relaunch JSAM within the same Safari user session. In order to re-launch JSAM, the user must exit Safari and then launch JSAM again.

Standard Application Support: MS Outlook

Remote users can use the Microsoft Outlook client on their PCs to access e-mail, their calendars, and other Outlook features through the system. This ability does not require changes to the Outlook client and does not require a network layer connection, such as VPN.

Refer to the *Supported Platforms Document* for details on operating system support and dependencies. See *Pulse Client-Side Changes Installation Reference* for details about registry changes made by JSAM.

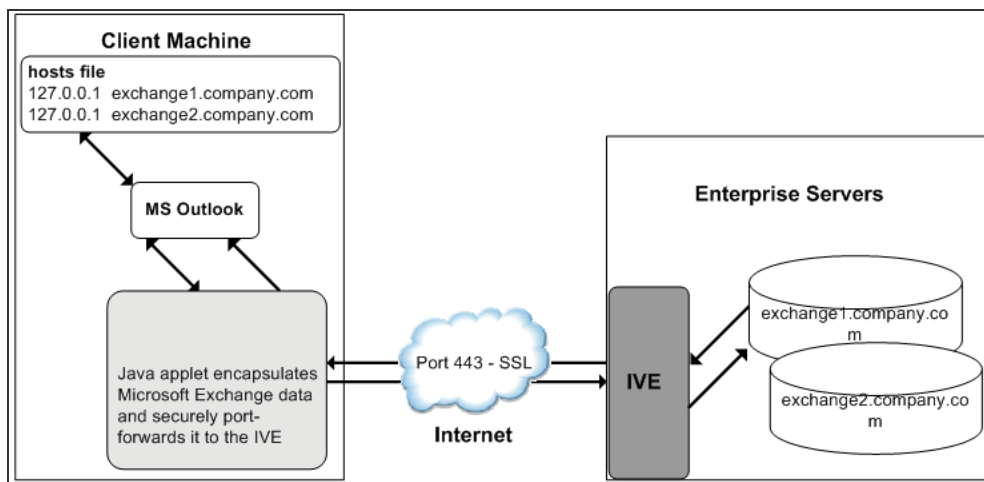
Also, note that the system does not support Outlook through SVW, since Outlook applications require HKLM registry key changes.

In order for this feature to work for remote users, the network settings of the user's PC must resolve the name of the Exchange Servers embedded in the Outlook client to the local PC (127.0.0.1, the default localhost IP address). We recommend that you configure the system to automatically resolve Exchange server hostnames to the localhost by temporarily updating the hosts file on a client computer through the automatic host-mapping option.

Client/Server Communication Using JSAM

The below figure describes the interactions between the Outlook client and an Exchange Server via Ivanti Connect Secure. The following figure assumes that the system is configured to perform automatic host-mapping.

The following figure depicts the Java Secure Application Manager and Enhanced MS Exchange Support:



1. The user starts the MS Outlook client. Outlook tries to contact the Exchange Server exchange1.yourcompany.com. The system resolves the Exchange Server hostname to 127.0.0.1 (localhost) through temporary changes to the hosts file.
2. Outlook connects to the Secure Application Manager running on the user's PC and then starts sending requests for e-mail.
3. The Secure Application Manager encapsulates and forwards all the requests from the Outlook client to the system over SSL.
4. The system unencapsulates the client data and looks in the MAPI request to find the target Exchange Server. The request is then forwarded to the target server.

5. Each request in the MAPI protocol encodes the target server for the request. When MAPI requests arrive from the Secure Application Manager, the system looks in each of them and dispatches them to the appropriate target server. This process works transparently even if there are multiple Exchange Servers.
6. The Exchange Server responds to the system with e-mail data.
7. The system encapsulates and forwards the response from the Exchange Server to the Secure Application Manager over SSL.
8. The Secure Application Manager unencapsulates the information sent from the system and forwards the normal MAPI response from the Exchange Server to the Outlook client.

Standard Application Support: Lotus Notes

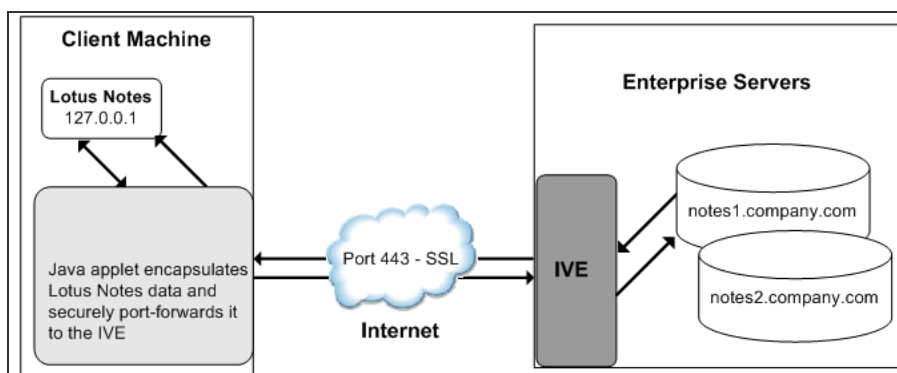
Remote users can use the Lotus Notes client on their PCs to access e-mail, their calendars, and other features through the system. This ability does not require a network layer connection, such as a VPN.

See the *Supported Platforms Document* for details on operating system support and dependencies.

Client/Server Communication Using JSAM

In order for this feature to work for remote users, they need to configure the Lotus Notes client to use "localhost" as their location setting (that is, their Home Location, Remote Location, or Travel Location setting). The Secure Application Manager then picks up connections requested by the Lotus Notes client. Figure 132 describes the interactions between the Lotus Notes client and a Lotus Notes Server via Ivanti Connect Secure.

The following figure depicts the Java Secure Application Manager and Enhanced Lotus Notes Support:



The above figure shows the Lotus Notes client location value to be configured to the localhost.

The user starts the Lotus Notes client with the location setting. The client uses the HTTP Tunnel proxy setting for its location setting. Note that you must set the HTTP Tunnel proxy setting to use localhost (or 127.0.0.1) as the proxy address and 1352 as the proxy port.

1. The Lotus Notes client connects to the Secure Application Manager and starts sending requests for e-mail.
2. The Secure Application Manager encapsulates and forwards requests from the Lotus Notes client to the system over SSL.
3. The system unencapsulates the client data and looks in the Lotus Notes request to find the target Lotus Notes Server. The request is then forwarded to the target server.

Each request in the Lotus Notes protocol encodes the target server for the request. When Lotus Notes requests arrive from the application proxy, the system obtains the target server information from the requests and dispatches the requests to the appropriate target server. Thus, this feature works transparently even if there are multiple Lotus Notes Servers accessed by a single user. Note that you must create JSAM ACLs on the system that enable access to these target servers.

4. The Lotus Notes Server responds with e-mail data to the system.
5. The system encapsulates and forwards the response from the Lotus Notes Server to the Secure Application Manager over SSL.
6. The Secure Application Manager unencapsulates the information sent from the system and forwards the normal response from the Lotus Notes Server to the Lotus Notes client.

Configuring the Lotus Notes Client

Before a remote user can connect from Lotus Notes to a Lotus Notes Server through the system, the user must edit the Lotus Notes client to set a Location document Proxy field to the PC's localhost port. The Location document edited should be the one used for remote access, such as the Remote Location or Travel Location setting. Setting the Proxy field to the PC's localhost port enables the system to connect to multiple Lotus Notes Servers, including those set up as pass-through servers.

You should use the following configuration in these cases:

- JSAM is configured to use Lotus Notes as a standard application.
- The Lotus Notes client can connect to multiple Lotus Notes servers.

To configure a Lotus Notes client for use with the system:

1. From the Lotus Notes client, choose **File > Mobile > Locations**.

2. Select the Location used for remote access and then click **Edit Location**.
3. In the Basics tab, click the **Proxy** icon.
4. In the Proxy Server Configuration box, enter **127.0.0.1:1352** in the HTTP Tunnel field.
5. Click **OK**.

Enabling Citrix Published Applications on the Citrix Native Client

When enabling Citrix published applications on the Citrix native client through the system, you must complete the following steps:

1. Specify custom application on JSAM to port forward.
2. Configure the Citrix metaframe server for published applications.
3. Configure the Citrix client for published applications.



- These instructions do not cover how to configure the standard Citrix application option. (For standard Citrix application instructions, use settings in the Users > Resource Profiles > Web > Web Applications/Pages page of the admin console.) You can enable both the standard Citrix application and the custom Citrix application-these settings do not impact each other.
- The system supports several mechanisms for intermediating traffic between a Citrix server and client, including the Terminal Services, JSAM, PSAM, VPN Tunneling, and hosted Java applets features.

Specifying Custom Applications on JSAM to Port Forward

When configuring JSAM to work with published applications, you must open two port-ports 80 and 1494. Each opened port creates a connection through JSAM to the Citrix Metaframe server.

To specify published applications for JSAM to port forward:

1. Add a custom application through JSAM. When adding the custom application, keep the following settings in mind:
 - Server name-For published applications, you must enter the Metaframe server's fully qualified DNS name, not its IP address.

- Server port-For published applications, enter 80 and 1494. (Create one entry for port 80 and another for port 1494.) If you have multiple Metaframe servers, you must configure all of them on the same ports.
 - Client port-For published applications, enter 80 and 1494. (Create one entry for port 80 and another for port 1494.)
2. If you have multiple internal domains, such as company-a.com and company-b.com, add DNS domains to the system using settings in the System > Network > Overview page of the admin console so that names such as app1.company-a.com and app2.company-b.com resolve correctly.
 3. If a remote user's PC is set up to use a Web proxy in Internet Explorer, configure the client machine to bypass the proxy server when the user launches applications that need to connect to the Secure Application Manager.
 4. Enable JSAM to associate IP loopback addresses with application servers on specific ports either by enabling JSAM to edit the hosts file on your users' systems or by creating an external DNS to route client application traffic to the JSAM applet.

Configuring the Citrix Metaframe Server for Published Applications

When enabling Citrix published applications through the system, you must enable the XML service DNS address resolution on the metaframe server. The following instructions describe how to do this on Metaframe XP.

To configure the Citrix metaframe server to work with the system:

1. Open the Citrix Management Console.
2. Right-click on the name of your server farm and click Properties.
3. Select the **MetaFrame Settings** tab.
4. Select the **Enable XML Service DNS address resolution** check box.
5. Click **OK**.

Configuring the Citrix Client for Published Applications

When enabling Citrix published applications through the system, you must create an ICA connection on each Citrix client using the instructions that follow.

To configure the Citrix client to work with the system:

1. Open the Citrix Program Neighborhood and choose the Add ICA Connection option.
2. In the Add New ICA Connection wizard, select the connection type that your computer uses to communicate.
3. In the next screen:
 1. Enter a description of the new ICA Connection.
 2. Select **TCP/IP + HTTP** as the network protocol.
 3. Select **Published Application**.
 4. Click **Server Location**, and then:
4. Deselect the **Use Default** check box.
 1. Click **Add** in the Locate Server or Published Application dialog box.
 2. Confirm that HTTP/HTTPS is selected from the Network Protocol list.
 3. Enter the metaframe server DNS in the Add Server Location Address dialog box.
 4. Enter 80 in the port field.
 5. Click **OK** in the Add Server Location Address dialog box and the Locate Server or Published Application dialog box.
 6. Select an application from the Published Application list.
5. Enter information in the remaining wizard screens as prompted.

Enabling Citrix Secure Gateways

When enabling Citrix secure gateways (CSGs) through the system, you must:

1. Disable Citrix NFuse as a standard application through the Users > Resource Profiles > Web > Web Applications/Pages page of the admin console.



You cannot enable the Citrix NFuse standard application and Citrix Secure Gateways (CSGs) custom applications through JSAM on the same device.

The system supports several mechanisms for intermediating traffic between a Citrix server and client, including the Terminal Services, JSAM, PSAM, VPN Tunneling, and hosted Java applets features.

2. Specify applications for JSAM to port forward by adding a custom application through JSAM. When adding the custom application, keep the following settings in mind:
 - Server name-For CSGs, you must enter the Citrix secure gateway server's fully qualified DNS name, not its IP address.
 - Server port-For CSGs, enter 443. If you have multiple Citrix secure gateway servers, you must configure all of them on the same port.
 - Client port-For CSGs, enter 443. (Create one entry for port 80 and another for port 443.)
3. If you have multiple internal domains, such as company-a.com and company-b.com, add DNS domains to the system using settings in the System > Network > Overview page of the admin console so that names such as app1.company-a.com and app2.company-b.com resolve correctly.
4. If a remote user's PC is set up to use a Web proxy in Internet Explorer, configure the client machine to bypass the proxy server when the user launches applications that need to connect to the Secure Application Manager.
5. Enable JSAM to associate IP loopback addresses with application servers on specific ports either by enabling JSAM to edit the hosts file on your users' systems or by creating an external DNS to route client application traffic to the JSAM applet.
6. Setup your Citrix Secure Gateway and confirm that it works on your desktop.
7. Add a bookmark to the end-users' home page that points to the list of Citrix secure gateway servers and use the Selective Rewrite feature to turn off rewriting for the URL.

Or, if you do not want to create a bookmark through the system, simply instruct users to access the URL using their Web browser's address bar instead of the system address bar.

Creating a JSAM Application Resource Profile

JSAM resource profiles configure JSAM to secure traffic to a client/server application. When you create a JSAM application resource profile, the JSAM client tunnels network traffic generated by the specified client applications to servers in your internal network.

When creating JSAM resource profiles, note that the resource profiles do not contain bookmarks. Therefore, end users will not see a link for the configured application in the end-user interface. To access the applications and servers that JSAM intermediates, users must first launch JSAM and then launch the specified application using standard methods (such as the Windows Start menu or a desktop icon).

Also note that when you enable JSAM or PSAM through rewriting autopolicies for Web resource profiles, the system automatically creates JSAM or PSAM autopolicies for you. You can only view these SAM policies through the appropriate Web resource profile-not through the SAM resource profile pages of the admin console.

To create a JSAM application resource profile:

1. In the admin console, choose **Users > Resource Profiles > SAM > Client Applications**.
2. Click **New Profile**.
3. From the Type list, choose **JSAM**.
4. From the Application list, select one of the following options.
 - **Custom**-Select this option to intermediate traffic to a custom application. Then:
 - In the Server name field, enter the name or IP address of the remote server. If you are using automatic host mapping, enter the server as it is known to the application. If you enter an IP address, note that end users must connect to JSAM using that IP address in order to connect to the specified server.
 - In the Server Port field, enter the port on which the remote server listens for client connections. For example, to forward Telnet traffic from a remote machine, specify port 23 for both the client port (on which JSAM listens) and the server port (on which the Telnet server listens).

To disable the registry change made by JSAM and restore the original copy of the etc/hosts file, users must uninstall the JSAM client using settings in the Preferences > Applications page of the end-user console. To re-enable the change, they need to reboot.

You can also use the restore system settings script. However, the restore system settings script cannot restore the hosts file successfully if you log in as a different user from the one that originally launched JSAM.

- In the Client Loopback IP field, provide a static loopback address. If you do not provide a static IP loopback address, the system assigns an IP loopback address dynamically.

When configuring an external DNS, do not use IP loopback addresses in the 127.0.2.x range because the system reserves IP loopback addresses in that range for use with Citrix NFuse.

If you want to modify a static loopback address for a JSAM application server configured on multiple ports, you must delete all applications referring to this application server and re-enter these applications with the new static loopback address.

- In the Client Port field, enter the port on which JSAM should listen for client application connections. Typically, the local port value is the same value as the server port; the local port value usually only differs for Linux or Macintosh non-root users who want to add applications for port forwarding that use ports under 1024.

You may configure more than one application on a single port, such as app1.mycompany.com, app2.mycompany.com, app3.mycompany.com. Either you assign a static loopback address or the system assigns a dynamic loopback address (127.0.1.10, 127.0.1.11, 127.0.1.12) to each application. JSAM then listens on these multiple loopback addresses on the specified port. For example, when there is traffic on 127.0.1.12 on the specified port, the system forwards the traffic to the app3.mycompany.com destination host.

5. Click **Add**.
 6. Select the **Allow JSAM to dynamically select an available port if the specified client port is in use** check box if JSAM is listening for multiple hosts on the same port and you want JSAM to select an available port when the client port you specify is taken. The client application must allow you to specify the port number for the connection in order to use this option.
 7. Select the **Create an access control policy allowing SAM access to these servers** check box to enable access to the list of servers specified in the Server column (enabled by default).
- **Lotus Notes** - Select this option to intermediate traffic from the Lotus Notes fat client application. Then, in the Autopolicy: SAM Access Control section, create a policy that allows or denies users access to the Lotus Notes server:
 - If it is not already enabled, select the **Autopolicy: SAM Access Control** check box.
 - In the **Resource** field, specify the application server to which this policy applies. You can specify the server as a fully-qualified hostname or an IP/netmask pair. For example, if the fully-qualified hostname is notes1.yourcompany.com, add notes1.yourcompany.com and notes1 to the Resource field.

- From the **Action** list, select **Allow** to enable access to the specified server or **Deny** to block access to the specified server.
- Click **Add**.



If you select the Lotus Notes option, or you configure the Lotus Notes client to connect to multiple Lotus Notes servers, you should configure the Lotus Notes client appropriately to work with Ivanti Connect Secure.

You can only use JSAM to configure access to one Lotus Notes application per user role.

- **Microsoft Outlook** - Select this option to intermediate traffic from the Microsoft Outlook application. Then:
 - Enter the hostname for each MS Exchange server in the Servers field. For example, if the fully-qualified hostname is exchange1.yourcompany.com, add exchange1.yourcompany.com to the Servers field.

You must enter the full name of the servers in this field since the system creates direct one-to-one mappings between the servers you enter here and IP addresses in the etc/hosts file. For more information about registry changes made by JSAM, see the *Client-side Changes Guide* on the Global Support Center.

The system does not support Outlook through SVW, since Outlook applications require HKLM registry key changes.

- Select the **Create an access control policy allowing SAM access to this server** check box to enable access to the server specified in the previous step (enabled by default).



You can only use JSAM to configure access to one Microsoft Outlook application per user role.

- **NetBIOS file browsing** - Select this option to tunnel NetBIOS traffic through JSAM. Then:
 - Enter the fully-qualified hostname for your application servers in the Servers field.

You must enter the full name of the servers in this field since the system creates direct one-to-one mappings between the servers you enter here and IP addresses in the etc/hosts file. For more information about registry changes made by JSAM, see the *Client-side Changes Guide* on the Global Support Center.

If you want to enable drive mapping on a Windows client machine, use the standard NetBIOS file browsing option. When you do, JSAM automatically modifies the registry to disable port 445 on Windows machines, which forces Windows to use port 137, 138, or 139 for drive-mapping. Windows users need to reboot one time to enable the registry change to take effect.

- Select the **Create an access control policy allowing SAM access to this server** check box to enable access to the server specified in the previous step (enabled by default).



You can only use JSAM to configure NetBIOS file browsing once per user role.

The system does not support NetBIOS file browsing through SVW, since NetBIOS requires HKLM registry key changes.

1. Enter a unique name and optionally a description for the resource profile. The system displays this information in the Client Application Sessions section of the end-user home page.
2. Click **Save** and **Continue**.
3. In the **Roles** tab, select the roles to which the resource profile applies and click **Add**.

The selected roles inherit the autopolicy created by the resource profile. If it is not already enabled, the system also automatically enables the SAM option in the Users > User Roles > *Role Name* > General > Overview page of the admin console for all of the roles you select.

4. Click **Save Changes**.

Specifying Applications for JSAM to Secure

Information in this section is provided for backwards compatibility. We recommend that you secure traffic using JSAM resource profiles instead, since they provide a simpler, more unified configuration method.

To specify applications for JSAM to secure:

1. In the admin console, choose **Users > User Roles > Select Role > SAM > Applications**.
2. Select **Add Application**.
3. Enter the name of the application and, optionally, a description. This information displays in the Client Application Sessions section of the end-user home page.
4. Choose either:

- **Standard application** - Select Citrix NFuse, Lotus Notes, or Microsoft Outlook/Exchange.

The system does not support the standard JSAM applications Outlook and Netbios file browsing through SVW, since these applications require registry key changes. However, the system does support the Citrix and Lotus Notes JSAM standard applications through SVW.

If you select the Lotus Notes option, or you configure the Lotus Notes client to connect to multiple Lotus Notes servers, you should configure the Lotus Notes client appropriately to work with Ivanti Connect Secure.

The system supports several mechanisms for intermediating traffic to the Lotus Notes, Microsoft Outlook, and Citrix applications.

- **Custom application** -

1. In the **Server name** field, enter the DNS name of the server or the server IP address. If entering the DNS name, enter name of the remote server as it is known to the application if you are using automatic host mapping.
2. Enter the server name.
3. In the **Server Port** field, enter the port on which the remote server listens for client connections.

For example, to forward Telnet traffic from a remote machine, specify port 23 for both the client port (on which JSAM listens) and the server port (on which the Telnet server listens).

To disable the registry change made by JSAM and restore the original copy of the etc/hosts file, users must uninstall the JSAM client using settings in the Preferences > Applications page of the end-user console. To re-enable the change, they need to reboot.

You can also use the restore system settings script. However, the restore system settings script cannot restore the hosts file successfully if you log in as a different user from the one that originally launched JSAM.

4. In the **Client Loopback IP** field, provide a static loopback address. If you do not provide a static IP loopback address, the system assigns an IP loopback address dynamically.

When configuring an external DNS, do not use IP loopback addresses in the 127.0.2.x range because the system reserves IP loopback addresses in that range for use with Citrix NFuse.

If you want to modify a static loopback address for a JSAM application server configured on multiple ports, you must delete all applications referring to this application server and re-enter these applications with the new static loopback address.

5. In the **Client Port** field, enter the port on which JSAM should listen for client application connections.

Typically, the local port value is the same value as the server port; the local port value usually only differs for Linux or Macintosh non-root users who want to add applications for port forwarding that use ports under 1024.

You may configure more than one application on a single port, such as app1.mycompany.com, app2.mycompany.com, app3.mycompany.com. Either you assign a static loopback address or the system assigns a dynamic loopback address (127.0.1.10, 127.0.1.11, 127.0.1.12) to each application. JSAM then listens on these multiple loopback addresses on the specified port. For example, when there is traffic on 127.0.1.12 on the specified port, the system forwards the traffic to the app3.mycompany.com destination host.

6. Select the **Allow Secure Application Manager to dynamically select an available port** ... check box if JSAM is listening for multiple hosts on the same port and you want JSAM to select an available port when the client port you specify is taken. The client application must allow you to specify the port number for the connection in order to use this option.
 7. Click **Add**.
5. If a remote user's PC is set up to use a Web proxy in Internet Explorer, configure the client machine to bypass the proxy server when the user launches applications that need to connect to the Secure Application Manager.

6. Add DNS domains to the system if you have multiple internal domains, such as company-a.com and company-b.com, so that names such as app1.company-a.com and app2.company-b.com resolve correctly:
 1. In the admin console, choose **System > Network > Overview**
 2. Under DNS Name Resolution, add a comma-separated list of domains in the to DNS Domains field.
 3. Click **Save Changes**.

Specifying Role Level JSAM Options

To specify JSAM options at the role level:

1. In the admin console, choose **Users > User Roles > Select Role > SAM > Options**.
2. Under Secure Application Manager options, select the options to enable for users:
 - **Auto-launch Secure Application Manager** - Select this option to automatically launches the Secure Application Manager when a user signs in. If you do not select this option, users must manually start the Secure Application Manager from the Client Applications Sessions section of the end-user home page.

Although you configure the Secure Application Manager to automatically launch when users sign into the system, users can override this setting through the Preferences > Applications page of the end-user console. If disabled from automatically launching, users need to manually start the Secure Application **Manager by clicking its link on the home page**.

- **Auto-uninstall Secure Application Manager** - Select this option to automatically uninstall the Secure Application Manager after users sign off.
- **Auto-allow application servers** - Select this option to automatically creates a SAM resource policy that allows access to the server specified in the PSAM application and server lists and the JSAM application list.

You may not see the Auto-allow option if you are using a new installation or if an administrator hides the option.

3. Under Java SAM Options, select the options to enable for users:
 - **User can add applications** - If enabled, users can add applications. For users to add applications, they need to know the application server DNS name and client/server ports.

When you enable this option, users can set up port forwarding to any host or port in your enterprise. Before providing users with the ability to add applications, please verify that this feature is consistent with your security practices. If a user adds an application, the application remains available to the user even if you later change disable the feature.

- **Automatic host-mapping** - If enabled, the Secure Application Manager edits the Windows PC hosts file and replaces entries of Windows application servers with localhost. These entries are changed back to the original data when a user closes the Secure Application Manager.

For the Java version of the Secure Application Manager to work, the client application needs to connect to the local PC on which the Secure Application Manager is running as the application server. The recommended process for mapping application servers to a user's local PC is to enable automatic host-mapping, which enables the system to automatically modify the PC's hosts file to point application servers to the PC's localhost for secure port forwarding. Alternatively, you can configure your external DNS server.

- **Skip web-proxy registry check** - If enabled, JSAM does not check a user's registry for a Web proxy. Some users do not have permissions to look at their registries, so if JSAM tries to look at their registries, then users see an error that they do not have permission. This option ensures that users do not see this message.
- **Auto-close JSAM window on sign-out** - If enabled, JSAM automatically closes when a user signs out of the device by clicking **Sign Out** in the browser window. JSAM continues to run if the user simply closes the browser window.

4. Click **Save Changes**.

Automatically Launching JSAM

Use the Launch JSAM tab to write a Web resource policy that specifies a URL for which the system automatically launches JSAM on the client. The system launches JSAM in two scenarios:

- When a user enters the URL in the Address field of the home page.
- When a user clicks a Web bookmark (configured by an administrator) on the home page to the URL.

This feature is useful if you enable applications that require JSAM but don't want to require users to run JSAM unnecessarily. This feature requires, however, that users access the URL through the home page. If users enter the URL in a browser Address field, the system does not serve the request.

The system provides tight integration with Citrix. If you specify Citrix as a standard JSAM application, the system automatically launches JSAM when a user selects an ICA file even if the URL is not configured as a resource policy.

To write a Launch JSAM resource policy:

1. In the admin console, choose **Users > Resource Policies > Web**.
2. If your administrator view is not already configured to show Launch JSAM policies, make the following modifications:
 - Click the **Customize** button in the upper right corner of the page.
 - Select the **Launch JSAM** check box.
 - Click **OK**.
3. Select the **Launch JSAM** tab.
4. On the JSAM Autolaunch Policies page, click **New Policy**.
5. Enter a name to label this policy (required) and a description of the policy (optional).
6. In the Resources section, specify the URLs to which this policy applies.



The resource policies configured for the JSAM auto launch policy must be a specific URL and not include wildcards. The URL should specify the entry point of the web application for which JSAM tunneling is needed.

7. In the Roles section, specify:
 - **Policy applies to ALL roles** - Choose this option to apply this policy to all users.
 - **Policy applies to SELECTED roles** - Choose this option to apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
 - **Policy applies to all roles OTHER THAN those selected below** - Choose this option to apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
8. In the Action section, specify:
 - **Launch JSAM for this URL** - The system downloads the Java Secure Application Manager to the client and then serves the requested URL.

- JSAM launches automatically for the specified URL only if a user enters the URL or selects a bookmark to the URL on the home page (Browsing > Bookmarks). The bookmark does not launch the application that is configured through JSAM, but launches JSAM itself.
 - **Don't Launch JSAM for this URL** - The system does not download the Java Secure Application Manager to the client for the requested URL. This option is useful if you want to temporarily disable JSAM auto-launching for the specified URLs.
 - **Use Detailed Rules** - To specify one or more detailed rules for this policy.
9. Click **Save Changes**.

Specifying Application Servers that Users Can Access

Information in this topic is provided for backwards compatibility. We recommend that you secure traffic using JSAM resource profiles instead, since they provide a simpler, more unified configuration method. Refer to the Specifying Application Servers that Users Can Access section in PSAM for more details.

Specifying Resource Level JSAM Options

Use the Options tab to specify the SAM resource option to match IP addresses to hostnames specified as resources in your SAM resource policies. When you enable this option, the system looks up IP addresses corresponding to each hostname specified in a SAM resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

When you enable this option, the system compiles a list of hostnames specified in the Resources field of each SAM resource policy. The system then applies the option to this comprehensive list of hostnames.



This option does not apply to hostnames that include wildcards and parameters.

To specify the SAM resource option:

1. In the admin console, choose **Users > Resource Policies > SAM > Options**.
2. Select IP based matching for Hostname based policy resources. When you select this option, the system looks up the IP address corresponding to each hostname specified in a Secure Application Manager resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

3. Click **Save Changes**.

Terminal Services

About Terminal Services

Use the Terminal Services feature to enable a terminal emulation session on a Windows terminal server, Citrix NFuse server, or Citrix Metaframe server. You can also use this feature to deliver the terminal services through the system, eliminating the need to use another Web server to host the clients.

The system supports several mechanisms for intermediating traffic between a Citrix server and client, including the Terminal Services, JSAM, PSAM, VPN Tunneling, and hosted Java applets features.

Terminal Services User Experience

From an end-user perspective, accessing secured terminal services resources through the system is simple. When you enable the Terminal Services feature for a user role, the end user simply needs to do the following tasks:

1. Specify the resource that the user wants to access-The user can specify the resource he wants to access by clicking a link or entering the resource in the system browse bar. Or, if you enable auto-launch for a bookmark, the system automatically launches the resource for the user when he signs into the device.
2. Enter credentials for the resource-When the user accesses a resource, the system prompts him to enter his username and password (if required by the resource). Or if you enable SSO, the system automatically sends this information to the resource without prompting the user. Once the resource verifies the credentials, the system launches the resource.

Users can access terminal services resources using the following methods:

- Session bookmarks-A session bookmark defines various information, including the server to which the user can connect, the terminal session's window parameters, and the username and password that the system sends to the Windows terminal server or Metaframe server. You can create any number of session bookmarks for a role, enabling the user to access multiple servers using different session bookmarks for each. (Users can simultaneously open multiple sessions to the same terminal server or to different servers.)
- URLs from other web sites-In most cases, users access session bookmarks directly from the end-user console. If you do not want to require users to sign into the end-user console to find and access terminal services links, you can create URLs on other web sites that point to session bookmarks that you have already created. Or, you can create URLs that include all of the parameters that you want to pass to the Terminal Services program, such as the host, ports, and terminal window parameters.



If you create links on external servers to terminal services bookmarks on the system and you are using multiple customized sign-in URLs, some restrictions occur.

- Ivanti Connect Secure browse bar-In addition to enabling users to link to terminal services links through bookmarks and URLs, you can also enable them to access these resources through the system browse bar on Windows systems. Users can access Citrix Metaframe or Nfuse servers by entering `ica://hostname` in the browse box. Or, users can access Microsoft terminal services or remote desktop sessions by entering `rdp://hostname` in the browse box.
- Server address-By entering a terminal server IP address or hostname, users can launch a remote desktop connection to any accessible server.

Task Summary: Configuring the Terminal Services Feature

To configure the Terminal Services feature:

1. Create resource profiles that enable access to Windows terminal servers or Citrix servers, include session bookmarks that link to those servers, and assign the session bookmarks to user roles using settings in the Users > Resource Profiles > Terminal Services page of the admin console.

We recommend that you use resource profiles to configure terminal services (as described here). However, if you do not want to use resource profiles, you can configure the Terminal Services feature using role and resource policy settings in the following pages of the admin console instead:

1. Create resource policies that enable access to Windows terminal servers and Citrix servers using settings in the Users > Resource Policies > Terminal Services > Access page of the admin console.
 2. Determine which user roles may access the Windows terminal servers and Citrix servers that you want to intermediate, and then enable Terminal Services access for those roles through the Users > User Roles > Select_Role > General > Overview page of the admin console.
 3. Create session bookmarks to your Windows terminal servers and Citrix servers using settings in the Users > User Roles > Select_Role > Terminal Services > Sessions page of the admin console.
2. (Optional.) Modify general role and resource options after configuring terminal services using resource profiles or roles and resource policies. Use the following pages of the admin console:

- (Optional.) Enable users to define their own terminal services sessions, specify the local devices to which users can connect, and set display and performance options using settings in the Users > User Roles > Select_Role > Terminal Services > Options page of the admin console. If you choose to enable users to define their own terminal services sessions, you must also create corresponding resource policies or resource profiles that enable access the specified resources, as explained in earlier in this topic.
 - (Optional.) Create links to a terminal services session that users can access from an external web site.
 - (Optional.) Enable the system to match IP addresses to hostnames using settings in the Users > Resource Policies > Terminal Services > Options page of the admin console.
3. (Citrix only) Specify where the system should obtain the Citrix client to upload to the users' systems through settings in the Users > User Roles > Select_Role > Terminal Services > Options page of the admin console.

Additionally, if you specify that the system should obtain a Citrix client from an external web site, you must:

- Create a Web access resource policy that enables access to the web site where the Citrix client resides through settings in the Users > Resource Policies > Web > Access > Web ACL page of the admin console.
- Create a Web caching resource policy through settings in the Users > Resource Policies > Web > Caching page of the admin console so the user's browser can deliver the Citrix client. (Note that you must select the Unchanged (do not add/modify caching headers) option.)

Terminal Services Execution

When a user tries to access a terminal services resource, the system completes the following steps to initiate and intermediate the terminal services session:

1. The system checks for a Java client.

To enable a terminal services session, the user either needs an RDP client on his system (to access a Windows terminal server) or an ICA client (to access a Citrix Metaframe server or server farm). These clients come in both Windows and Java versions and enable the user to run an application on the server while only transmitting keyboard, mouse, and display information over the network.

The system enables you to upload a Java version of the RDP or ICA client through a terminal services resource profile (but not role). If you have uploaded a client to the system and specified that the system always use it to run your users' terminal sessions, the system launches the specified Java client.

-
2. (Citrix only.) If necessary, the system checks for a Windows client.

If you have not uploaded a Java client, the system checks for a Windows version of the ICA client. If it cannot find a Windows ICA client, it installs the version you specified in the Users > User Roles > Role > Terminal Services > Options page of the admin console.

-
-
3. The system checks for the terminal services proxy.

To intermediate a Windows or Citrix session, the user either needs a Terminal Services proxy on his system or a Citrix Services Client proxy. The system checks for the appropriate proxy on the user's computer, and if it cannot find it, installs a new one. Depending on the user's rights, the system either uses an ActiveX component or Java component to install the proxy.

-
-
-
4. The proxy tries to invoke the Windows client.

Once the system has confirmed that a proxy is installed on the user's computer, the proxy attempts to invoke the Windows RDP or ICA client. If successful, the client initiates the user's terminal services session and the proxy intermediates the session traffic.

-
-
-
-
5. The proxy tries to invoke the Java client.

If a Winitiates the user's terminal services session and the proxy intermediates the session traffic.

For informatdows client is not present on the user's machine (for instance, because it was deleted or because the user does not have the proper privileges to install it), but you have uploaded one to the system through the terminal services resource profile, the system uses the uploaded Java applet to launch the session.

As part of the installation, the system asks the user if he wants to always use the Java client or only for this session. The system then stores the user's preference as a persistent cookie. Once the Java client is installed, the client inion about the specific files installed by the system when you enable the Terminal Services feature, as well as the rights required to install and run the associated clients, see the *Client-side Changes Guide* on the Global Support Center.

Configuring Citrix to Support ICA Load Balancing

The Service Terminal Services feature supports connecting to Citrix server farms in which published applications are preconfigured (as described later in this topic). The system does not support load balancing configurations in which Nfuse servers dynamically retrieve a list of Citrix published applications within a server farm.

Citrix Load Balancing Overview

The system supports the following Citrix load balancing scenario:

1. The Citrix administrator makes a published application available to multiple Citrix servers in a farm by generating a custom ICA file. The generated ICA file contains a parameter called HTTPBrowserAddress that points to the IP address and port number of the master browser (that is, the server that performs the load balancing).
2. When the ICA client attempts to launch a published application on the user's computer, it uses the HTTPBrowserAddress parameter to connect to the master browser.
3. The master browser pings servers in the farm to determine their respective loads and returns the IP address of the least busy server to the ICA client.

4. The ICA client uses the IP address returned by the master browser to connect to the appropriate terminal server.

Configuring Citrix Load Balancing

For the system to work properly with a Citrix farm, you must configure the Citrix farm and Ivanti Connect Secure as described in the following steps. Note that these instructions are based on using a Citrix Metaframe Presentation Server 3.0.

1. On the Citrix server, enable a server (or multiple servers) in your farm as a master browser:
 1. Right-click a server in the Metaframe Farm and select **Properties**.
 2. Select **Metaframe Settings**.
 3. Enter the TCP/IP port for the Citrix XML service.
2. On the Citrix server, publish the applications that are hosted on MetaFrame XP servers in the farm:
 1. Right-click the Applications link and select **Publish applications**.
 2. Specify which desktop or application to publish.
 3. Follow the prompts in the wizard.
 4. Specify the list of servers that host the application you are publishing and click Finish.
3. The specified published application appears in the server farm.
4. On the Citrix server, generate a corresponding Citrix ICA file for the published application:
 1. Select the application you published in Step 2 and select **Create ICA file**.
 2. Follow the prompts in the wizard.

3. On the TCP/IP + HTTP Server page, enter the name of the HTTP browser server and the port number. (The port should match the Citrix XML Service port that you set up in Step 1).
4. Save the ICA file.
5. On Ivanti Connect Secure, upload the ICA file using settings in either of the following admin console pages:
 - Users > User Roles > *Role* > Terminal Services > Sessions
 - Users > Resource Profiles > *Profile*
6. On Ivanti Connect Secure, create a resource policy for the HTTP browser server and port entered in Step 3.
7. On Ivanti Connect Secure, test the configuration by launching the bookmark as an end user.



One of the Citrix servers in the farm performs the load balancing, not Ivanti Connect Secure. If the ICA client is already installed on the user's desktop then administrator rights are not required.

For more information about the rights required to use the Terminal Services feature, see *Client-Side Changes Installation Reference*.


If the XML response from the master browser contains the hostname, it will not work through Ivanti Connect Secure. To ensure that the response is in dot-port format (an IP address), clear the Enable XML service DNS address resolution check box during the browser server configuration. This option controls whether the destination Citrix server is represented as a hostname or as an IP address.

About Terminal Services Resource Profiles

Terminal Services resource profile configuration instructions vary depending on whether you want to configure access to a Windows terminal server (which requires an RDP client) or Citrix terminal server (which requires an ICA client). Furthermore, if you choose to configure access to a Citrix server using a custom ICA file, you include many of your configuration settings in the ICA file itself and therefore do not need to configure them through the system. If you configure access to a Citrix server using the default ICA file on the system, however, you must configure additional settings.

You may want to create multiple bookmarks for the same terminal services resource in order to provide easy access to multiple applications. For instance, the server defined in your resource profile may provide access to multiple applications (such as Siebel and Outlook). To easily provide access to each of these applications, you can create resource profile bookmarks to each. Or, you may want to use multiple bookmarks to configure single sign-on to one application, but not another.

When configuring session bookmarks, note that:

- To change the host or ports for a terminal services session bookmark created through a resource profile, you must edit the values through the resource profile's Resource tab (not its Bookmark tab).
- You can only assign session bookmarks to roles that you have already associated with the resource profile—not all of the roles defined on the system. To change the list of roles associated with the resource profile, use settings in its Roles tab.
-  Session bookmarks simply control which links to display to users—not which resources the users can access. For example, if you enable access to a terminal server through a resource profile but do not create a corresponding session bookmark to that server, the user can still access the server by entering it into the Address box of the home page.
- Make sure to enter a unique set of parameters when defining a terminal services bookmark. If you create two bookmarks that contain the same set of parameters, the system deletes one of the bookmarks from the end-user view. You can still see both bookmarks, however, in the administrator console.

Configuring a Windows Terminal Services Resource Profile

This topic describes how to configure a terminal services resource profile that enables access to a Windows terminal server using an RDP client.

Users can use RDP7 features through the Terminal Services if an RDP7 client is present. However, the true multi-monitor and bidirectional audio features of RDP7 are not supported with this release.

To create a Windows terminal services resource profile:

1. In the admin console, select **Users > Resource Profiles > Terminal Services**.
2. Click **New Profile**. Or select an existing profile from the list.

3. Select **Windows Terminal Services** from the Type list.
4. Enter a unique name and optionally a description for the resource profile. (This name becomes the default session bookmark's name.)
5. Specify the server and port to which this resource profile should connect in the Host field. When entering the server, you may enter a hostname or IP address.
6. Enter the port on which the terminal server listens in the Server port box. (By default, the system populates this box with port number 3389.)
7. Select the **Create an access control policy allowing Terminal Service access to this server** check box to enable access to the server specified in the Server Port box (enabled by default).
8. If you want to enable intermediation using a Java client, select **Enable Java support** and then specify which Java client the system should use.
9. Click **Save** and **Continue**.
10. Select the roles to which the resource profile applies in the Roles tab and click Add.

The selected roles inherit the autopolicy and session bookmarks created by the resource profile. If it is not already enabled, the system also automatically enables the Terminal Services option in the Users > User Roles > Select_Role > General > Overview page of the admin console for all of the roles you select.

11. Click **Save Changes**.
12. (Optional.) Modify the default session bookmark created by the system in the Bookmarks tab and/or create new ones. By default, the system creates a session bookmark to the server defined in the Host box and displays it to all users assigned to the role specified in the Roles tab.)

Defining a Hosted Java Applet Autopolicy

Hosted Java applet autopolicies enable you to store terminal services Java clients directly on the system without employing a separate Web server to host them. You can then associate these Java applets with the resource profile and specify that the system always use them to intermediate traffic, or that the system fall back to the applets when other terminal services clients are not available on the user's system.

Although you can use a Java applet to intermediate traffic to an SSO-enabled resource, we do not recommend it because the applet may require the user's password to be presented as plain text.

A default Premier Java RDP applet is shipped with each device and cannot be deleted. The HOB applet is available through the New Terminal Services Resource Profile window and the Users > User Roles > Users > Terminal Services > Options window. To use the Ivanti supplied HOB applet, you must contact Ivanti Support to purchase a license including the number of concurrent users you want to support.

The HOB applet is similar to any other Java applet accessed through the system or uploaded to the system. You must install a code-signing certificate to avoid seeing a warning similar to "This applet was signed by "Ivanti Connect Secure" but Java cannot verify the authenticity of the signature's certificate. Do you trust this certificate?" Install a valid Applet signing certificate (JavaSoft) in the Configuration > Certificates > Code-signing Certificates window.



The HOB applet is for RDP connections and appears only for Windows Terminal Services. It is not applicable for Citrix Terminal Services profiles. The supported HOB version is 4.1.0794.

You can purchase HOB applets directly from HOB; however, Ivanti will support them only to the extent of uploading them. If you have any problems configuring or running the applet, you must contact HOB support.

To create a hosted Java applet autopolicy:

1. Create a terminal services resource profile.
2. Select **Enable Java support** within the resource profile.

3. Select the Java applet that you want to associate with the resource profile from the Applet to use list. Or, if the applet that you want to use is not currently available in the list, click Edit Applet. Then:

- Click **New Applet** to add an applet to this list. Or, select an existing applet and click Replace (to replace an existing applet with a new applet) or **Delete** (to remove an applet from the system).



If you replace an existing archive, make sure that the new applet archive contains all of the necessary files for the applet to successfully launch and run. If the associated HTML for the applet refers to files that do not exist in the new archive, then the applet will not function correctly.

The system only allows you to delete applets that are not currently in use by a Web or terminal services resource profile.

If you select the Enable Java support option and have a custom ICA file that you uploaded to the system, your HTML file is auto-populated with references to your custom ICA file. No additional HTML code needs to be added.

1. Enter a name to identify the applet in the Name box. (This applies to new and replaced applets only.)
2. Browse to the applet that you want to upload to the system. You can upload applets (.jar or .cab files) or archives (.zip, .jar, and .tar files) that contain applets and all of the resources that the applets need. (This applies to new and replaced applets only.)
3. If the file that you selected is an archive that contains the applet, select the Uncompress jar/cab file check box. (This applies to new and replaced applets only.)
4. Click **OK** and **Close Window**.



When you select an applet in the Java Applets dialog box, you are loading third-party software onto the Ivanti product. By clicking OK, you are agreeing to the following terms on behalf of yourself (as purchaser of the equipment) or the organization that purchased the Ivanti product, as applicable:

By loading third party software onto the Ivanti product, you are responsible for obtaining all rights necessary for using, copying, and/or distributing such software in or with the Ivanti product. Ivanti is not responsible for any liability arising from use of such third party software and will not provide support for such software. The use of third-party software may interfere with the proper operation of the Ivanti product and/or Ivanti software, and may void any warranty for the Ivanti product and/or Ivanti software.

4. Create an HTML page definition in the HTML box that includes references to your Java applets. The maximum size of the HTML that can be specified is 25k. Then, fill in any required attributes and parameters.

If you are using HTML generated by the system, make sure to search the HTML code for "`__PLEASE_SPECIFY__`" and update the code as necessary.



If you select Hob-Ivanti RDP Applet from the Applet to Use menu, you must select the Configure HTML for the default applet check box in order to edit the HTML. Otherwise, the default HTML is used. By default, the proxy mode is disabled in the Hob - RDP Applet.

To enable the proxy mode, add the following:

```
<parameter name="proxymode" value="http">
```

If your proxy requires authentication, add the following to the Hob-RDP Applet:

```
<parameter name="proxyuser" value="<username>">
```

```
<parameter name="proxypassword" value="<password>">
```


You can also add any additional HTML or JavaScript that you choose to this Web page definition. The system rewrites all of the code that you enter in this box.

Make sure to enter unique HTML in this box. If you create two bookmarks with the same HTML code, the system deletes one of the bookmarks in the end-user view. You can still see both bookmarks, however, in the administrator console.

For dynamic drive mapping to work with HOB Applet 4.1.0794, you must enable both the AUTOLDM and **TWAUTOMAPDRIVE** parameters. See the *Premier Java Applet Configuration Options document* located on the Ivanti support site for more details on these two parameters.

5. Select **Use this Java applet as a fallback mechanism** to use this applet only when the Windows client fails to launch. Or select **Always use this Java applet** to use this applet regardless of whether or not the Windows client launches.
6. Click **Save Changes**.

Defining a Bookmark for a Windows Terminal Services Profile

When you create a terminal services resource profile, the system automatically creates a bookmark that links to the terminal server that you specified in the resource profile. The system allows you to modify this bookmark as well as create additional bookmarks to the same terminal server.

To configure resource profile bookmarks for Windows terminal services:

1. In the admin console, select **Users > Resource Profiles > Terminal Services > Resource Profile Name > Bookmarks**.
2. Click the appropriate link in the Bookmark column if you want to modify an existing session bookmark. Or, click **New Bookmark** to create an additional session bookmark.

Although it is generally easiest to create a resource profile session bookmark through the resource profile configuration page, you can choose to create one through the user roles page as well.

3. (Optional.) Change the name and description of the session bookmark. (By default, the system populates and names the session bookmark using the resource profile name.)
4. Specify how the terminal emulation window should appear to the user during a terminal session by configuring options in the Settings area of the bookmark configuration page.
5. Pass user credentials from the system to the terminal server so that users can sign onto the terminal server without having to manually enter their credentials. You can do this by configuring options in the Session area of the bookmark configuration page.
6. Allow users to access specific applications on the terminal server by configuring options in the Start Application area of the bookmark configuration page. In addition, you can use settings in this area to define auto-launch and session reliability options.
7. Allow users to access local resources such as printers and drives through the terminal session by configuring options in the Connect Devices area of the bookmark configuration page.
8. Specify how the terminal emulation window should appear to the user during a terminal session by configuring options in the Desktop Settings area.
9. Specify the roles to which you want to display the session bookmarks if you are configuring the session bookmark through the resource profile pages, under Roles:
 - **ALL selected roles**-Displays the session bookmark to all of the roles associated with the resource profile.
 - **Subset of selected roles**-Displays the session bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL Selected Roles list and click Add to move them to the Subset of selected roles list.
10. Click **Save Changes**.

Creating a Windows Terminal Services Bookmark Through the User Roles Page

It is generally easiest to create a terminal services bookmark through the resource profile configuration pages. However, you can choose to create a resource profile session bookmark through the user roles page using the following instructions:

1. In the admin console, select **Users > User Roles > Select Role > Terminal Services > Sessions**.
2. Click **Add Session**.
3. Select **Terminal Services Resource Profile** from the Type list. (This option displays only after you have already created a terminal services resource profile.)
4. Select an existing resource profile that connects to a Windows terminal server on the system. (The system automatically populates the Host and Server Port boxes using settings from the selected resource profile.)
5. Click **OK**. (If you have not already associated the selected role with the resource profile, the system automatically makes the association for you. The system also enables any access control policies for the role that are required by the resource profile.)
6. If this role is not already associated with the selected resource profile, the system displays an informational message. If you see this message, click **Save Changes** to add this role to the resource profile's list of roles and to update the profile's autopolicies as required. Then, repeat the previous steps to create the session bookmark.

When you create a resource profile session bookmark through the user roles page (instead of the standard resource profiles page), the system only associates the generated session bookmark with the selected role. The system does not assign the session bookmark to all of the roles associated with the selected resource profile.

7. (Optional.) Change the name and description of the session bookmark. By default, the resource profile name is used as the bookmark name.

8. Configure the bookmark's settings.

Defining Display Options for the Windows Terminal Services Session

When configuring a terminal services bookmark, you can specify how the terminal emulation window should appear to users during their terminal sessions.

To define display options and auto-launch options:

1. Create a terminal services bookmark or edit an existing bookmark.
2. Scroll to the Settings area of the bookmark configuration page.
3. Select an option from the Screen Size drop-down list if you want to change the size of the terminal services window on the user's workstation. The default window size is full screen.



If you select the Full Screen option and are connecting to a Windows terminal server, the system modifies the user's hosts file to display the correct hostname in the terminal services window. If the user does not have the proper rights to modify the hosts file, the system displays the loopback address instead.

Also note that to restore the hosts file to its original state after running the terminal services window, the user must properly close his application. Otherwise, other applications that use the hosts file (such as JSAM and Host Checker) might not run properly. The user can also restore his hosts file to its original state by rebooting his system or by renaming the backup hosts file (**hosts_ive.bak**).

4. Select **8-bit, 15-bit, 16-bit, 24-bit, or 32-bit** color from the Color Depth list if you want to change the color-depth of the terminal session data. The default color depth is 8-bit.
5. Click **Save Changes**.

Defining SSO Options for the Windows Terminal Services Session

When configuring a terminal services bookmark, you can configure the system to pass user credentials from the system to the terminal server so that the user does not have to manually enter his username and password. The system passes the specified credentials when a user clicks the session bookmark. If the credentials fail, the server prompts the user to manually enter his username and password.

To define single sign-on options:

1. Create a terminal services bookmark or edit an existing bookmark.
2. Scroll to the Authentication area of the bookmark configuration page.
3. Specify the username to pass to the terminal server. You can enter a static username or a variable. Enter the <username> variable to pass the username stored in the system's primary authentication server. Or use the following syntax to submit the username for the secondary authentication server: <username@SecondaryServerName> or <username[2]>.
4. Select Password if you want to specify a static password or select Variable Password if you want to use the password stored in the system's primary or secondary authentication server. To use the password from the primary authentication server, enter the <password> variable. Or use the following syntax to submit the password for the secondary authentication server: <Password@SecondaryServerName> or <Password[2]>.
5. Click **Save Changes**.

Defining Application Settings for the Windows Terminal Services Session

When configuring a terminal services bookmark, you can specify that users can only access specific applications on the terminal server.

To define applications that users can access:

1. Create a terminal services bookmark or edit an existing bookmark.

2. Scroll to the Start Application area of the bookmark configuration page.
3. Select the Launch seamless window check box to have the Windows application server manage the display of the application. This allows an application's windows to behave in the same way as an application running on a Windows application server, regardless of the user's desktop environment.



If SSO is not configured, seamless window is supported only on Remote Desktop Protocol (RDP) 6.0 and later.

The Launch seamless window check box is applicable only for servers running Windows 2008 and later.

Enter the server alias name (applicable only for servers running Windows 2008 and later) in the Alias name box.

1. Specify where the application's executable file resides on the terminal server in the Path to application box (visible only when you clear Launch seamless window). For example, you might enter the following directory for the Microsoft Word application:

C:\Program Files\Microsoft Office\Office10\WinWord.exe

2. Specify where the terminal server should place working files for the application in the Working directory box. For example, you might specify that Microsoft Word should save files to the following directory by default:

C:\Documents and Settings\username\My Documents



You can use session variables such as <username> and <password> in the Path to application and Working directory boxes. For example, when specifying an application path, you might want to include the <username> variable to personalize the location. For example:
C:\Documents and Settings\<username>\My Documents.

3. Select the **Auto-launch** check box if you want to automatically launch this Terminal Service session bookmark when users sign into device. When you select this option, the system launches the terminal services application in a separate window after the user signs in.
4. Click **Save Changes**.

Defining Device Connections for the Windows Terminal Services Session

When configuring a terminal services bookmark, you can specify local resources that users can access through the terminal session.



The system does not support providing users access to local resources when intermediating traffic using Java applets. Therefore, if you select the Enable Java Applets option when creating a Windows Terminal Services resource profile, note that the Connect Devices options described below might not work.

When you enable local resources through the terminal server, each user can only access his own local resources. For instance, user 1 cannot see user 2's local directories.

To define local resources that users can access:

1. Create a terminal services bookmark or edit an existing bookmark.
2. Scroll to the Connect Devices area of the bookmark configuration page.
3. Select **Connect local drives** to connect the user's local drive to the terminal server, enabling the user to copy information from the terminal server to his local client directories.
4. Select **Connect local printers** to connect the user's local printers to the terminal server, enabling the user to print information from the terminal server to his local printer.
5. Select **Connect COM Ports** to connect the user's COM ports to the terminal server, allowing communication between the terminal server and the devices on his serial ports.

6. Select **Allow Clipboard Sharing** to allow the contents of the clipboard to be shared between the user's host computer and the terminal server. Because of limitations in RDP client earlier than version 6.0, clearing the Allow Clipboard Sharing option will automatically disable the Connect local drives, Connect local printers, and Connect COM Ports options.
7. Select **Connect smart cards** to allow users to use smart cards to authenticate their remote desktop sessions.



Smart cards are supported by Microsoft Remote Desktop Protocol versions 5.1 and later.

8. Select **Sound Options to enable sound** during the remote session. Select Bring to this computer to redirect audio to the local computer. Select Leave at remote computer to play the audio only at the server.



Sound options are supported by Microsoft Remote Desktop Protocol versions 5.1 and later.

9. Select **Use Multiple Monitors** to support multiple monitors connected to the client computer during the remote session.



Multiple monitors are supported by Microsoft Remote Desktop Protocol versions 8.1 and later.

10. Select the **Network Level Authentication** check box to enable the NLA at the bookmark level.
11. Select the **Allow Smartcard with Network Level Authentication** check box to enable smart cards and NLA simultaneously.



This option is applicable to non-cross-domain certificates.

12. Select the **Use Remote Microphones** check box to support microphones connected to the client computer during the remote session.
13. Click **Save Changes**.

Defining Desktop Settings for the Windows Terminal Services Session

When configuring a terminal services bookmark, you can specify how the terminal emulation window should appear to the user during a terminal session.



The options in this topic only apply to Windows Terminal Services bookmarks.

To define display settings for the users' sessions:

1. Create a terminal services bookmark or edit an existing bookmark
2. Scroll to the **Display Settings** area of the bookmark configuration page.
3. Select **Desktop background** to display a wallpaper background to users. If you do not select this option, the background is blank.
4. Select **Show contents of window** while dragging to show the contents of the Windows Explorer window while users move the windows on their desktops.
5. Select **Menu and window animation** to animate the movement of windows, menus, and lists.
6. Select **Themes to allow users** to set Windows themes in their terminal server windows.
7. Select **Bitmap Caching** to improve performance by minimizing the amount of display information that is passed over a connection.
8. Select **Font Smoothing (RDP 6.0 onwards)** to make text smoother and easier to read. This option only works on Windows Vista computers running RDP clients that are version 6.0 or later.
9. Select **Desktop Composition (RDP 6.0 onwards)** to allow desktop composition. With desktop composition, individual windows no longer draw directly to the screen. Instead, their drawing is redirected to video memory, which is then rendered into a desktop image and presented on the display.

10. Click **Save Changes**.

Creating a Citrix Terminal Services Resource Profile Using Default ICA Settings

This topic describes how to configure access to a Citrix Metaframe server using a default ICA configuration file.

To create a Citrix Terminal Services resource profile that uses default ICA settings:

1. In the admin console, select **Users > Resource Profiles > Terminal Services**.
2. Click **New Profile**. Or select an existing profile from the list.
3. Select **Citrix using default ICA** from the Type list.
4. (Existing resource profiles only) If you want to customize the default ICA file that comes with the system, click the Open link, customize the file, and upload it.
5. Enter a unique name and optionally a description for the resource profile. (This name becomes the default session bookmark's name.)
6. Specify the server and port to which this resource profile should connect in the Host box. When entering the server, you may enter a hostname or IP address.
7. Enter the port on which the terminal server listens in the Server Port field. (By default, the system populates this field with port number 1494 for Citrix.)
8. Select the **Create an access control policy allowing Terminal Service access to this server** check box to enable access to the server specified in the Server Port box (enabled by default).
9. Enable intermediation using a Java client by selecting Enable Java support and then specifying which Java client the system should use.
10. Click **Save** and **Continue**.

11. Select the roles to which the resource profile applies in the Roles tab and click Add.

The selected roles inherit the autopolicy and session bookmarks created by the resource profile. If it is not already enabled, the system also automatically enables the Terminal Services option in the Users > User Roles > Select_Role > General > Overview page of the admin console for all of the roles you select.

12. Click **Save Changes**.
13. (Optional.) Modify the default session bookmark created by the system in the Bookmarks tab and/or create new ones. (By default, the system creates a session bookmark to the server defined in the Host box and displays it to all users assigned to the role specified in the Roles tab.)

Defining a Bookmark for a Citrix Profile Using Default ICA Settings

When you create a Terminal Services resource profile, the system automatically creates a bookmark that links to the terminal server that you specified in the resource profile. The system enables you to modify this bookmark as well as create additional bookmarks to the same terminal server.

To configure resource profile bookmarks for Citrix Terminal Services using default ICA settings:

1. In the admin console, select Users > Resource Profiles > Terminal Services > Select Resource Profile > Bookmarks.
2. Click the appropriate link in the Bookmark column if you want to modify an existing session bookmark. Or, click New Bookmark to create an additional session bookmark.



Although it is generally easiest to create a resource profile session bookmark through the resource profile configuration page, you can choose to create one through the user roles page as well.

3. (Optional.) Change the name and description of the session bookmark. (By default, the system populates and names the session bookmark using the resource profile name.)
4. Specify how the terminal emulation window should appear to the user during a terminal session use configuration options in the Settings area of the bookmark configuration page.
5. Pass user credentials from the system to the terminal server so users can sign onto the terminal server without having to manually enter their credentials. You can do this by using the configuration options in the Session area of the bookmark configuration page.
6. Allow users to access specific applications on the terminal server by using configuration options in the Start Application area of the bookmark configuration page. In addition, you can use settings in this section to define auto-launch and session reliability options.
7. Allow users to access local resources such as printers and drives through the terminal session by using the configuration options in the Connect Devices section of the bookmark configuration page.
8. Specify the roles to which you want to display the session bookmark if you are configuring the session bookmark through the resource profile pages, under Roles:
 - **ALL selected roles**-Displays the session bookmark to all of the roles associated with the resource profile.
 - **Subset of selected roles**-Displays the session bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL selected roles list and click Add to move them to the Subset of selected roles list.
9. Click **Save Changes**.

Creating a Citrix Terminal Services Bookmark Through the User Roles Page

It is generally easiest to create a terminal services bookmark through the resource profile configuration pages, as explained in the previous topic. However, you can choose to create a resource profile session bookmark through the user roles page using the following instructions:

1. In the admin console, select **Users > User Roles > Select_Role > Terminal Services> Sessions**.
2. Click **Add Session**.
3. Choose **Terminal Services Resource Profile** from the Type list. (The system does not display this option if you have not already created a terminal services resource profile.)
4. Select an existing resource profile that connects to a Citrix server using the default ICA file. (The system automatically populates the Host and Server Port fields using settings from the selected resource profile.)
5. Click **OK**. (If you have not already associated the selected role with the resource profile, the system automatically makes the association for you. The system also enables any access control policies for the role that are required by the resource profile.)
6. If this role is not already associated with the selected resource profile, the system displays an informational message. If you see this message, click **Save Changes** to add this role to the resource profile's list of roles and to update the profile's autopolicies as required. Then, repeat the previous steps to create the session bookmark.



When you create a resource profile session bookmark through the user roles page (instead of the standard resource profiles page), the system only associates the generated session bookmark with the selected role. The system does not assign the session bookmark to all of the roles associated with the selected resource profile.

7. (Optional.) Change the name and description of the session bookmark. By default, the resource profile name is used as the session bookmark name.


8. Configure the bookmark's settings.

Defining Display Options for the Citrix Terminal Services Session

When configuring a terminal services bookmark, you can specify how the terminal emulation window should appear to users during their terminal sessions.

To define display, auto-launch, and session reliability options:

1. Create a terminal services bookmark or edit an existing bookmark.
2. Scroll to the Settings area of the bookmark configuration page.
3. Select an option from the Screen Size drop-down list if you want to change the size of the terminal services window on the user's workstation. The default window size is full screen.
4. Select **8-bit, 15-bit, 16-bit, 24-bit, or 32-bit** color from the Color Depth list if you want to change the color-depth of the terminal session data. The default color depth is 8-bit.

 When configuring a Citrix session bookmark, note that the setting you choose here and the user's local desktop setting both affect the client's color-depth display. If these settings do not match, the user sees the lower of the two color-depths during his session. For example, if you select 16-bit color during configuration, but the user's local desktop is set to 8-bit, the user sees 8-bit color depth during his session.

5. Click **Save Changes**.

Defining SSO Options for the Citrix Terminal Services Session

When configuring a terminal services bookmark, you can configure the system to pass user credentials from the system to the terminal server so that the user does not have to manually enter his username and password. The system passes the specified credentials when a user clicks the session bookmark. If the credentials fail, the server prompts the user to manually enter his username and password.

To define single sign-on options:

1. Create a terminal services bookmark or edit an existing bookmark.
2. Scroll to the Authentication area of the bookmark configuration page.
3. Specify the username to pass to the terminal server in the Username field. You can enter a static username or a variable. Enter the <username> variable to pass the username stored in the system's primary authentication server. Or use the following syntax to submit the username for the secondary authentication server:
<username@SecondaryServerName> or <username[2]>.
4. Select Password if you want to specify a static password or select Variable Password if you want to use the password stored in the system's primary or secondary authentication server. To use the password from the primary authentication server, enter the <password> variable. Or use the following syntax to submit the password for the secondary authentication server:
<Password@SecondaryServerName> or <Password[2]>.
5. (Default ICA file and listed applications only.) Select Use domain credentials to pass the user's cached domain credentials to the Citrix Metaframe server (also called pass-through authentication). When you select this option, the system uses the Citrix Program Neighborhood client to intermediate the Citrix terminal session.



If you want to download the Program Neighborhood client, select Users > User Roles > Select_Role > Terminal Services > Options in the admin console and enter the URL in the Download from URL box. See the Citrix web site for the location of the latest Program Neighborhood client cab file.

When you select the Use domain credentials option, you must also enable SSO through the user's settings file (appsrv.ini). If the user has already successfully signed into the Metaframe server using cached domain credentials, this setting should already be enabled. Otherwise, you or the user must:

- **Set EnableSSOnThruICAFile=On** in appsrv.ini. You can locate appsrv.ini in the %HOMEPATH%\Application Data\ICAClient directory.
 - **Set UseLocalUserAndPassword=On** in the ICA file.
6. If you have not enabled SSO through the INI file, the user is prompted to manually enter his credentials when he tries to access the Metaframe server through the system.
 7. Click **Save Changes**.

Defining Application, Auto-Launch, and Session Reliability Settings for the Citrix Terminal Services Session

When configuring a terminal services bookmark, you can specify that users can only access specific applications on the terminal server.

To define applications that users can access:

1. Create a terminal services bookmark or edit an existing bookmark.
2. Scroll to the Start Application area of the bookmark configuration page.

3. Select the Launch seamless window check box to have the Windows application server manage the display of the application. This allows an application's windows to behave in the same way as an application running on a Windows application server, regardless of the user's desktop environment.



If SSO is not configured, seamless window is supported only on Remote Desktop Protocol (RDP) 6.0 and later.

Enter the server alias name in the Alias Name field (applicable only for servers running Windows 2008 and later).

4. Specify where the application's executable file resides on the terminal server in the Path to application box (visible only when you clear Launch seamless window). For example, you might enter the following directory for the Microsoft Word application:

C:\Program Files\Microsoft Office\Office10\WinWord.exe

5. Specify where the terminal server should place working files for the application in the Working directory field. For example, you might specify that Microsoft Word should save files to the following directory by default:

C:\Documents and Settings\\My Documents



You can use system session variables such as <username> and <password> in the Path to application and Working directory boxes. For example, when specifying an application path, you might want to include the <username> variable to personalize the location. For example, C:\Documents and Settings\\My Documents.

6. Select the **Auto-launch** check box if you want to automatically launch this terminal service session bookmark when users sign into the device. When you select this option, the system launches the terminal services application in a separate window when the user signs in.

7. Select **Session Reliability and Auto-client reconnect** to keep ICA sessions active and on the user's screen when network connectivity is interrupted. Users continue to see the application they are using until the network connectivity resumes or the session reliability time-out has expired (the time-out value is defined by the Citrix product). Enter the port to use in the Port to be enabled box.
8. Click **Save Changes**.

Defining Device Connections for the Citrix Terminal Services Session

When configuring a terminal services bookmark, you can specify local resources that users can access through the terminal session.

For the Connect Devices settings to take effect, they must also be enabled on the Metaframe server. For example, if you enable Connect Drives on the system, but disable it on the Metaframe server, then the Metaframe server will block access to local drives. Note that if you clear the Configure access to local resources check box, the settings on the Metaframe server take effect.

To define local resources that users can access:

1. Create a terminal services bookmark or edit an existing bookmark.
2. Scroll to the Connect Devices area of the bookmark configuration page.
3. Select **Connect local drives** to connect the user's local drive to the terminal server, enabling the user to copy information from the terminal server to his local client directories.
4. Select **Connect local printers** to connect the user's local printers to the terminal server, enabling the user to print information from the terminal server to his local printer.
5. Select **Connect COM Ports** to connect the user's COM ports to the terminal server, allowing communication between the terminal server and the devices on his serial ports.

6. When you enable local resources through the terminal server, each user can only access his own local resources. For instance, user 1 cannot see user 2's local directories.
7. Click **Save Changes**.

Creating a Citrix Resource Profile That Uses a Custom ICA File

Use this type of resource profile to enable a terminal session to a Citrix Metaframe server using settings that you specify in a customized ICA file. Use custom ICA files to enable terminal sessions to Citrix Metaframe servers or NFuse servers governing Citrix server farms (in other words, to load balance). You may also use custom ICA files to link to single servers, if necessary. When you select this option, the system uses the session parameters defined in the specified custom ICA file.

To enable the connection between the system and the Citrix server farm, you must use the TCP/IP+HTTP protocol for browsing and specify the Citrix Metaframe or NFuse server port and IP address. The system does not support UDP port-forwarding.

To create a Citrix resource profile that uses a custom ICA file:

1. In the admin console, select **Users > Resource Profiles > Terminal Services**.
2. Click **New Profile**. Or select an existing profile from the list.
3. Select **Citrix using custom ICA file** from the Type list.
4. Specify the ICA file that contains the session parameters that you want use in the Custom ICA File box. Note that you may download and customize the following ICA files from the system:

- ICA file that comes with the system-To customize this file, click the Open link, save the file to your local machine, customize the file as required, and upload it back to the system using the Browse option. If you customize this file, you must replace the following parameters in the default.ica file: <CITRIX_CLIENT_NAME>, <APPDATA> and <TARGET_SERVER>.
- ICA file that you have already associated with the resource profile-To customize this file, click the Current ICA File link, save the file to your local machine, and customize the file as required. Once you make changes, you must upload the revised version using the Browse option.

Before uploading the ICA file, you should test it to make sure it initiates the Citrix session correctly. To test, create an ICA file and access it directly. If the file displays the Citrix session correctly then it should work through the system.

If SSO is configured in the custom ICA bookmark, seamless mode is ignored and the application is launched in non-seamless mode.

5. Enter a unique name and optionally a description for the resource profile. (This name becomes the default session bookmark's name.)

6. Enable access to the servers specified in the custom ICA file:

1. Select the **Autopolicy: Terminal Services Access Control** check box.
2. Specify the Metaframe servers to which you want to enable access in the Resource field.
3. Choose **Allow** to enable access to the specified resource or **Deny** to block access to the specified resource from the Action list.
4. Click **Add**.

7. Enable intermediation using a Java client by selecting Enable Java support.

If you select the **Enable Java support** option and have a custom ICA file that you uploaded to the system, your HTML file is auto-populated with references to your custom ICA file. No additional HTML code needs to be added.

8. Click **Save** and **Continue**.

9. Select the roles to which the resource profile applies in the Roles box and click **Add**.

The selected roles inherit the autopolicy and session bookmarks created by the resource profile. If it is not already enabled, the system also automatically enables the Terminal Services option in the Users > User Roles > Select_Role > General > Overview page of the admin console for all of the roles you select.

10. Click **Save Changes**.

11. (Optional) Modify the default session bookmark created by the system in the Bookmarks tab and/or create new ones. (By default, the system creates a session bookmark to the server defined in your custom ICA file and displays it to all users assigned to the role specified in the Roles tab.)

Defining a Bookmark for a Citrix Profile Using a Custom ICA File

When you create a terminal services resource profile, the system automatically creates a bookmark that links to the terminal server that you specified in the resource profile. You can modify this bookmark as well as create additional bookmarks to the same terminal server.

To configure resource profile bookmarks for Citrix Terminal Services using custom ICA settings:

1. In the admin console, select **Users > Resource Profiles > Terminal Services > Select_Resource_Profile > Bookmarks**.

Click the appropriate link in the Bookmark column if you want to modify an existing session bookmark. Or, click **New Bookmark** to create an additional session bookmark.

Although it is generally easiest to create a resource profile session bookmark through the resource profile configuration page, you can choose to create one through the user roles page as well.

2. (Optional.) Change the name and description of the session bookmark. (By default, the system populates and names the session bookmark using the resource profile name.)
3. Pass user credentials from the system to the terminal server so that users can sign onto the terminal server without having to manually enter their credentials. You can do this by configuring options in the Session area of the bookmark configuration page.
4. Automatically launch this terminal service session bookmark when a user signs in to the device by selecting the Auto-launch check box. When you select this option, the system launches the terminal services application in a separate window when the user signs in.

5. Under Roles, specify the roles to which you want to display the session bookmark:
 - **ALL selected roles**-Displays the session bookmark to all of the roles associated with the resource profile.
 - **Subset of selected roles**-Displays the session bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL selected roles list and click Add to move them to the Subset of selected roles list.
6. Click **Save Changes**.

Creating a Citrix Profile That Lists Published Applications

Citrix created published applications to satisfy the need for security. It is dangerous to allow any executable to be run on the server. With published applications, only applications that are allowed to be run are published.

These published applications are displayed on the system index page as terminal services bookmarks.



The Citrix Desktop Toolbar Viewer is enabled only for XenDesktop. It is not enabled for XenApp. If you require the Citrix Desktop Toolbar Viewer, use the XenDesktop configuration on Connect Secure. Do not configure a desktop as part of the Citrix Listed Applications feature.

To create a Citrix profile that lists published applications:

1. In the admin console, select **Users > Resource Profiles > Terminal Services**.
2. Click **New Profile**.
3. Select **Citrix Listed Applications** from the Type list.
4. Enter a unique name and optionally a description for the resource profile. This name becomes the default session bookmark's name.
5. Enter the IP address and port of the Citrix MetaFrame server where the XML service is running.

You do not need to enter the port number if you are using the default value. The default port is 80 (if SSL is selected, the default port is 443).

You can enter more than one server. If the connection fails on one server, the next server in the list is used.

6. Click the **Use SSL for connecting to Citrix XML Service** check box to send the password through SSL instead of cleartext.



Although cleartext is supported, we recommend you always use SSL to avoid any security issues.

7. Enter the username, password, and domain name for connecting to the Citrix Metaframe server where the XML service is running.

You can enter variable credentials such as <USERNAME> and <PASSWORD>. If you use variable credentials, the Subset of selected Applications option is disabled in the Bookmarks window.

When the user accesses the application list, their credentials are submitted to the Citrix XML service, substituting the session context variables <USERNAME> and <PASSWORD>. Only the user's specific applications (as determined by the Citrix administrator) are returned.

8. Enable access to the servers specified in the custom ICA file:
 1. Select the **Autopolicy: Terminal Services Access Control** check box.
 2. Specify the Metaframe servers to which you want to enable access in the Resource field.
 3. Choose **Allow** to enable access to the specified resource or **Deny** to block access to the specified resource from the Action list.
 4. Click **Add**.
9. Enable intermediation using a Java client by selecting **Enable Java support** and then specifying which Java client to use.

10. Click **Save** and **Continue**.
11. Select the roles to which the resource profile applies in the Roles tab and click Add.

The selected roles inherit the autopolicy and session bookmarks created by the resource profile. If it is not already enabled, the system also automatically enables the Terminal Services option in the Users > User Roles > Select_Role > General > Overview page of the admin console for all of the roles you select.

12. Click **Save Changes**.
13. (Optional.) Modify the default session bookmark created by the system in the Bookmarks box and/or create new ones.

Defining a Bookmark for a Citrix Profile Listing Applications

When you create a terminal services resource profile, the system automatically creates a bookmark that links to the terminal server that you specified in the resource profile. You can modify this bookmark as well as create additional bookmarks to the same terminal server.

To configure resource profile bookmarks for Citrix terminal services list applications:

1. In the admin console, select **Users > Resource Profiles > Terminal Services > Resource_Profile > Bookmarks**.

Click the appropriate link in the Bookmark column if you want to modify an existing session bookmark. Or, click New Bookmark to create an additional session bookmark.

Although it is generally easiest to create a resource profile session bookmark through the resource profile configuration page, you can choose to create one through the user roles page as well.

2. (Optional.) Change the name and description of the session bookmark. By default, the resource profile name is used as the session bookmark name.

3. Under Applications, select the applications you want available to the end user.
 - **ALL Applications**-Allow all executables on the server to be available to the end user.
 - **Subset of selected applications**-Select the executables from the Available list and click Add to allow only those applications to be run. The Available list is automatically populated from the Metaframe server.

This option is disabled when you enter variable credentials, such as <USERNAME> and <PASSWORD> while defining the resource profile.
4. Under Settings, specify how the terminal emulation window should appear to users during their terminal sessions.



You cannot change the IP address or XML Service running port for connecting to the XML Service or the Java client to use for intermediation.

- Select an option from the Screen Size drop-down list if you want to change the size of the terminal services window on the user's workstation.
 - (Optional.) Select **8-bit, 15-bit, 16-bit, 24-bit, or 32-bit** color from the Color Depth list if you want to change the color-depth of the terminal session data.
5. Under Session, you can configure the system to pass user credentials from the system to the terminal server so that the user does not have to manually enter his username and password.

1. Specify the username to pass to the terminal server in the Username box. You can enter a static username or a variable.
2. Select **Password** if you want to specify a static password or select Variable Password if you want to use the password stored in the system's primary or secondary authentication server.
3. Select **Use domain credentials** to pass the user's cached domain credentials to the Citrix Metaframe server (also called pass-through authentication). When you select this option the system uses the Citrix Program Neighborhood client to intermediate the Citrix terminal session.



If you want to download the Citrix Program Neighborhood client, select **Users > User Roles > Role Name > Terminal Services > Options** of the admin console and enter the following URL in the Download from URL box:
`http://download2.citrix.com/FILES/en/products/client/ica/client9230/wficat.cab`

When you select the Use domain credentials option, you must also enable SSO through the user's settings file (appsrv.ini).

6. Under Connect Devices, specify which user devices to connect to the terminal server.
 - **Connect local drives**-Connect the user's local drive to the terminal server, enabling the user to copy information from the terminal server to his local client directories.
 - **Select Connect local printers**-Connect the user's local printers to the terminal server, enabling the user to print information from the terminal server to his local printer.

- **Select Connect COM Ports**-Connect the user's COM ports to the terminal server, allowing communication between the terminal server and the devices on his serial ports.
7. Under Roles, specify the roles to which you want to display the session bookmark:
 8. Click **Save Changes**.

Creating Session Bookmarks to Your Terminal Server

When you enable the Terminal Services option through the admin console, you can create session bookmarks to your terminal server. A terminal services session bookmark defines information about the terminal server to which users can connect and (optionally) applications that they can use on the terminal server. The session bookmarks that you define appear on the Terminal Services panel in the end-user console for users who map to the appropriate role.

You can use two different methods to create terminal services session bookmarks:

- Create session bookmarks through existing resource profiles (recommended)-When you select this method, the system automatically populates the session bookmark with key parameters (such as the session type) using settings from the resource profile. Additionally, while you are creating the associated resource profile, the system guides you through the process of creating any required policies to enable access to the session bookmark.
- Create standard session bookmarks-With this option, you must manually enter all session bookmark parameters during configuration. Additionally, you must enable access to the Terminal Services feature and create resource policies that enable access to the servers defined in the session bookmark.



If you enable the Terminal Services option but do not give users the ability to create their own session bookmarks, make sure that you configure session bookmarks for them. Otherwise, users cannot use this feature.

You can also enable users to create their own session bookmarks on the homepage and browse to the terminal servers using the system browse bar. Or, you can create links from external sites to a terminal services bookmark.

Creating Advanced Terminal Services Session Bookmarks

The information in this topic is provided for backwards compatibility. We recommend that you configure access to Windows terminal servers and Citrix servers through resource profiles instead, because they provide a simpler, more unified configuration method. Resource profile also contain features (such as the ability to use Java RDP clients to support Macintosh and Linux users) which are not available through roles.

Make sure to enter a unique set of parameters when defining a terminal services bookmark. If you create two bookmarks that contain the same set of parameters, the system deletes one of the bookmarks from the end user view. You can still see both bookmarks in the administrator console.

To create a session bookmark for terminal sessions:

1. In the admin console, select **Users > User Roles > Role > Terminal Services > Sessions**.
2. Click **Add Session**.
3. Select **Standard** in the Type drop-down list.
4. Specify the type of user session you want to create from the Session Type list:
 - **Windows Terminal Services**-Enables a terminal session to a Windows terminal server.
 - **Citrix using default ICA**-Enables a terminal services session to a Citrix Metaframe server. When you select this option, the system uses the default Citrix session parameters.

(Existing sessions only.) You can also use the Open link to open the system's default ICA file, which you can then save to your local machine and customize as required. If you customize this file, you must replace the following parameters in the default.ica file: <CITRIX_CLIENT_NAME>, <APPDATA>, and <TARGET_SERVER>.

- **Citrix using custom ICA file**-Enables a terminal services session to a Citrix Metaframe or NFuse server governing a Citrix server farm. When you select this option, the system uses the session parameters defined in the specified custom ICA file, thus removing the Session Reliability, Start Application, and Connect Devices configuration items from the current page.



Because the system does not support UDP port-forwarding, you must use the TCP/IP+HTTP protocol for browsing and specify the Citrix Metaframe or NFuse server port and IP address to enable the connection between Connect Secure and the Citrix server farm.

5. Enter a name and (optionally) a description for the session bookmark.
6. In the Host field, specify the hostname or IP address of the Windows terminal server or Metaframe terminal server.
7. In the Client Port and Server Port fields, enter the ports on which the user client communicates and terminal server listens.

If you specify a client port and the terminal services client is unable to bind to this port, then the terminal services client will fail. However, if you leave the Client Port field blank, the terminal services, Citrix Services Client dynamically selects an available port.

8. (Windows Terminal Services and Citrix using default ICA only) If you want to specify the screen size and color depth options for the terminal emulation window, use configuration options in the Settings section.

9. If you want to pass user credentials from the system to the terminal server, enabling users to sign onto the terminal server without having to manually enter their credentials, use configuration options in the Session section.
10. If you only want to allow users to access specific applications on the terminal server, use configuration options in the Start Application section of the bookmark configuration page. In addition, you can use settings in this section to define auto-launch and session reliability options.
11. (Windows Terminal Services and Citrix using default ICA only) If you want to allow users to access local resources such as printers and drives through the terminal session, use configuration options in the Connect Devices section of the bookmark configuration page.
12. (Windows Terminal Services only) If you want to specify how the terminal emulation window should appear to the user during a terminal session, use configuration options in the Desktop Settings section.
13. Click **Save Changes** or **Save + New**.

Defining Screen Size and Color Depth Options for the Terminal Services Session

When configuring a terminal services bookmark, you can specify how the terminal emulation window should appear to users during their terminal sessions.

The options in this section only apply to Windows Terminal Services bookmarks, Citrix using default ICA bookmarks and Citrix listed applications bookmarks.

To define display, auto-launch, and session reliability options:

1. Create a terminal services session bookmark or edit an existing session bookmark.
2. Scroll to the Settings section of the bookmark configuration page.

3. Select an option from the Screen Size drop-down list if you want to change the size of the terminal services window on the user's workstation. The default window size is full screen.

If you select the Full Screen option and are connecting to a Windows terminal server, the system modifies the user's hosts file in order to display the correct hostname in the terminal services window. If the user does not have the proper rights to modify the hosts file, the system displays the loopback address instead.

Also note that in order to restore the hosts file to its original state after running the terminal services window, the user must properly close his application. Otherwise, other applications that use the hosts file (such as JSAM and Host Checker) might not run properly. The user can also restore his hosts file to its original state by rebooting his system or by renaming the backup hosts file (hosts_ive.bak).

4. Select a value from the Color Depth list if you want to change the color-depth of the terminal session data. The default color depth is 8-bit.

When configuring a Citrix session bookmark, note that the setting you choose here and the user's local desktop setting both affect the client's color-depth display. If these settings do not match, the user sees the lower of the two color-depths during his session. For example, if you choose 16-bit color during system configuration, but the user's local desktop is set to 8-bit, the user sees 8-bit color depth during his session.

5. Click **Save Changes** or **Save + New**.

Defining SSO Options for the Terminal Services Session

When configuring a terminal services bookmark, you can pass user credentials from the system to the terminal server so that the user does not have to manually enter his username and password. The system passes the specified credentials when a user clicks the session bookmark. If the credentials fail, the server prompts the user to manually enter his username and password.

To define single sign-on options:

1. Create a terminal services session bookmark or edit an existing session bookmark.
2. Scroll to the Authentication section of the bookmark configuration page.
3. In the Username field, specify the username to pass to the terminal server. You can enter a static username or a variable. Enter the <username> variable to pass the username stored in the system's primary authentication server. Or use the following syntax to submit the username for the secondary authentication server:
<username@SecondaryServerName> or <username[2]>.
4. Select **Password** if you want to specify a static password or select Variable Password if you want to use the password stored in the system's primary or secondary authentication server. To use the password from the primary authentication server, enter the <password> variable. Or use the following syntax to submit the password for the secondary authentication server:
<Password@SecondaryServerName> or <Password[2]>.
5. (Citrix using default ICA or listed applications) Select **Use domain credentials** to pass the user's cached domain credentials to the Citrix Metaframe server (also called pass-through authentication). When you select this option, the system uses the Citrix Program Neighborhood client to intermediate the Citrix terminal session.



If you want to download the Program Neighborhood client, go to the Users > User Roles > Select Role > Terminal Services > Options page of the admin console and enter the following URL in the Download from URL field:

<https://downloadplugins.citrix.com/Windows/CitrixReceiver.exe>

When you select the Use domain credentials option, you must also enable SSO through the user's settings file (appsrv.ini). If the user has already successfully signed into the Metaframe server using cached domain credentials, this setting should already be enabled. Otherwise, you or the user must:

- Set EnableSSOnThruICAFile=On in appsrv.ini. You can locate appsrv.ini in the %HOMEPATH%\Application Data\ICAClient directory.
- Set UseLocalUserAndPassword=On in the ICA file.

If you have not enabled SSO through the INI file, the user is prompted to manually enter his credentials when he tries to access the Metaframe server through the system.

6. Click **Save Changes** or **Save + New**.

Defining Application Settings for the Terminal Services Session

When configuring a terminal services bookmark, you can specify that users can only access specific applications on the terminal server. Additionally, you can define auto-launch and session reliability options for the session.

To define applications that users can access:

1. Create a terminal services session bookmark or edit an existing session bookmark.
2. Scroll to the Start Application section of the bookmark configuration page.

If you specify Citrix using custom ICA file in the Session Type configuration section, the Start Application configuration item is not available.

3. (Windows Terminal Services and Citrix using default ICA only) In the Path to application field, specify where the application's executable file resides on the terminal server. For example, you might enter the following directory for the Microsoft Word application:

C:\Program Files\Microsoft Office\Office10\WinWord.exe

4. (Windows Terminal Services and Citrix using default ICA only) In the Working directory field, specify where the terminal server should place working files for the application. For example, you might specify that Microsoft Word should save files to the following directory by default:

C:\Documents and Settings\\My Documents

You can use session variables such as <username> and <password> in the Path to application and Working directory fields. For example, when specifying an application path, you might want to include the <username> variable to personalize the location. For example: C:\Documents and Settings\\My Documents.

5. (Citrix using default ICA only) Select **Session Reliability and Auto-client reconnect to keep ICA sessions active** and on the user's screen when network connectivity is interrupted. Users continue to see the application they are using until the network connectivity resumes or the session reliability time-out has expired (the time-out value is defined by the Citrix product). Enter the port to use in Port to be enabled.
6. Select the Auto-launch check box if you want to automatically launch this Terminal Service session bookmark when users sign into the device. When you select this option, the system launches the terminal services application in a separate window when the user signs in.
7. Click **Save Changes** or **Save + New**.

Defining Device Connections for the Terminal Services Session

When configuring a terminal services bookmark, you can specify local resources that users can access through the terminal session.

The options in this section only apply to Windows Terminal Services bookmarks and Citrix using default ICA bookmarks.

The Connect Devices options that you specify at the role-level control whether end users can enable or disable access to local resources when they configure their own bookmarks. These role-level options do not control whether users can access local resources through a bookmark created by the system administrator.

To define local resources that users can access:

1. Create a terminal services session bookmark or edit an existing session bookmark.

2. Scroll to the **Connect Devices** section of the bookmark configuration page.

If you specify Citrix using custom ICA file in the Session Type configuration section, the Connect Devices configuration item is not available.

3. Select **Connect local drives** to connect the user's local drive to the terminal server, enabling the user to copy information from the terminal server to his local client directories.
4. Select **Connect local printers** to connect the user's local printers to the terminal server, enabling the user to print information from the terminal server to his local printer.
5. Select **Connect COM Ports** to connect the user's COM ports to the terminal server, allowing communication between the terminal server and the devices on his serial ports.
6. (Windows Terminal Services only) Select **Allow Clipboard Sharing** if you want to allow the contents of the clipboard to be shared between the user's host computer and the terminal server. Due to the limitations in the pre-6.0 versions of the RDP client, disabling the Allow Clipboard Sharing option will automatically disable the Connect local drives, Connect local printers, and Connect COM Ports options.

When you enable local resources through the terminal server, each user can only access his own local resources. For instance, user 1 cannot see user 2's local directories.

7. (Windows Terminal Services only) Select **Connect smart cards** to allow users to use smart cards to authenticate their remote desktop sessions.
8. (Windows Terminal Services only) Select **Sound Options** to enable sound during the remote session. Choose Bring to this computer to redirect audio to the local computer. Choose Leave at remote computer to play the audio only at the server.



Smart cards and sound options are supported by **Microsoft Remote Desktop Protocol versions 5.1** and later.

9. (Windows Terminal Services only) Select **Use Multiple Monitors** to support multiple monitors connected to the client computer during the remote session.



Multiple monitors are supported by Microsoft Remote Desktop Protocol versions 8.1 and later.

10. (Windows Terminal Services only) Select the **Use Remote Microphones** check box to support microphones connected to the client computer during the remote session.
11. Select the **Network Level Authentication** check box to enable the NLA at the bookmark level.
12. Select the **Allow Smartcard with Network Level Authentication** check box to enable smart cards and NLA simultaneously.



This option is applicable to non-cross-domain certificates.

13. Click **Save Changes** or **Save + New**.

Defining Desktop Settings for the Terminal Services Session

When configuring a terminal services bookmark, you can specify how the terminal emulation window should appear to the user during a terminal session.

The options in this section only apply to Windows Terminal Services bookmarks.

To define display settings for the users' sessions:

1. Create a terminal services session bookmark or edit an existing session bookmark.
2. Scroll to the Display Settings section of the bookmark configuration page.

3. Select **Desktop background** if you want to display a wallpaper background to users. If you do not select this option, the background is blank.
4. Select **Show contents of window while dragging** if you want to show the contents of the Windows Explorer window while users move the windows on their desktops.
5. Select **Menu and window animation** if you want to animate the movement of windows, menus, and lists.
6. Select **Themes** if you want to allow users to set Windows themes in their terminal server windows.
7. Select **Bitmap Caching** if you want to improve performance by minimizing the amount of display information that is passed over a connection.
8. Select **Font Smoothing (RDP 6.0 onwards)** to make text smoother and easier to read. This option only works on Windows Vista computers running RDP clients that are version 6.0 or later.
9. Select **Desktop Composition (RDP 6.0 onwards)** to allow desktop composition. With desktop composition, individual windows no longer draw directly to the screen. Instead, their drawing is redirected to video memory, which is then rendered into a desktop image and presented on the display.
10. Click **Save Changes** or **Save + New**.

Creating Links from an External Site to a Terminal Services Session Bookmark

When creating a link to a terminal services session bookmark from another site, you can construct either of the following types of URLs:

- URL that includes all necessary parameters-Create a URL that includes all of the parameters that you want to pass to the terminal services program, such as the host, ports, and terminal window parameters.

When constructing the URL, use the following syntax:

```
https://<SASeriesAppliance>/dana/term/winlaunchterm.cgi?<param1>=<value1>&<param2>=<value2>
```

When constructing your URL, you can use the case-insensitive parameter names described in Table 92. If you want to include more than one parameter in the session bookmark, string them together using ampersand characters (&). For example:

```
https://YourSA.com/dana/term/winlaunchterm.cgi?host=yourtermserver.yourdomain.com&type=Windows&clientPort=1094&serverPort=3389&user=john&password=abc123&screenSize=fullscreen
```

URL to terminal services bookmark-Create a URL that simply points to a session bookmark that you have already created on the system. When constructing the URL, use the following syntax:

```
https://<SASeriesAppliance>/dana/term/winlaunchterm.cgi?bmname=<bookmarkName>
```

Within the URL, only define the bmName parameter.

When using the system to host Terminal Services session bookmarks, you must:

- Enable the User can add sessions option in the Users > User Roles > Select Role > Terminal Services > Options page. If you do not select this option, users cannot link to the Terminal Services session bookmarks from external sites.
- Create a policy that prevents the system from rewriting the link and the page that contains the link using settings in the Users > Resource Policies > Web > Rewriting > Selective Rewriting page of the admin console.

Additionally, we recommend that you use https protocol instead of http. Otherwise, when users launch the session bookmark, they see an insecure site warning.



If you create links on external servers to Terminal Services bookmarks on the system and you are using multiple customized sign-in URLs, some restrictions occur.

The following table describes Case-Insensitive Terminal Services Session Bookmark Parameter Names:

Parameter Name	Description and Possible Values	Example
host	Required. Hostname or IP address of the Windows terminal server or Metaframe terminal server.	https://YourSystem.com/dana/term/winlaunchterm.cgi?bmname=<bookmarkname>
type	Type of terminal server. Possible values include Windows or Citrix.	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&type=Windows
clientPort	Port on which the user client communicates.	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&clientPort=1094
serverPort	Port on which the terminal server listens. Default values are 3389 for Windows and 1494 for Citrix.	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&serverPort=3389
user	Username to pass to the terminal server. You can enter a static username, such as JDoe, or a variable username such as <user> or <username>	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&user=jDoe

Parameter Name	Description and Possible Values	Example
password	Password to pass the terminal server.	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&user=jDoe&password= <password>
bookmarkname	Specifies the session bookmark name	https://YourSystem.com/dana/term/winlaunchterm.cgi?bookmarkname= <bookmarkname>
screenSize	Terminal services window's size. Possible values:fullScreen 800x6001024x76 81280x1024	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&screenSize=fullScreen
colorDepth	Terminal services window's color depth, in bits. Possible values:815162432	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&colorDepth=32
startApp	Specifies the path of an application executable to start.	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&startApp=C:\Program Files\Microsoft Office\Office10\WinWord.exe
startDir	Specifies where the terminal server should place working files for the application.	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&startapp=C:\temp

Parameter Name	Description and Possible Values	Example
connect Drives	Specifies whether the user can connect his local drive to the terminal server. Possible values:YesNo	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&connectDrives=Yes
connect Printers	Specifies whether the user can connect his local printer to the terminal server. Possible values:YesNo	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&connectPrinters=Yes
connect ComPorts	Specifies whether the user can connect his COM ports to the terminal server. Possible values:YesNo	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&connectComPorts=Yes
allowclipboard	Specifies whether the user can share the contents of the clipboard between the user's host computer and the terminal server. Possible values:YesNo	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&allowclipboard=Yes

Parameter Name	Description and Possible Values	Example
desktopbackground	Specifies whether to display your current wallpaper setting. Possible values:YesNo	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&desktopbackground=Yes
showDragContents	Specifies whether to show the contents of the Windows Explorer window while moving the window around your desktop. Possible values:YesNo	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&showDragContents=Yes
showMenuAnimation	Specifies whether to animate the movement of windows, menus, and lists. Possible values:YesNo	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&showMenuAnimation=Yes
themes	Specifies whether to allow users to set Windows themes in their terminal server windows. Possible values:YesNo	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&themes=Yes

Parameter Name	Description and Possible Values	Example
bitmap caching	Specifies whether to improve performance by minimizing the amount of display information that must be passed over a connection. Possible values:YesNo	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&bitmapcaching=Yes
fontsmoothing	Specifies whether to make text smoother and easier to read. This option only works on Windows Vista computers running RDP clients that are version 6.0 or later. Possible values:YesNo	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&fontsmoothing=Yes
desktopcomposition	Specifies whether to enable desktop composition. Possible values:YesNo	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&desktopcomposition=Yes

Parameter Name	Description and Possible Values	Example
soundoptions	Specifies whether to enable sound. Possible values:0-disable sound1-bring sound to this computer2-leave sound at remote computer	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&soundOptions=1
multiMon	Specifies whether to allow users to use all the monitors for the remote session.Possible values:YesNo	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&multiMon=Yes

Specifying General Terminal Services Options

Users can create their own terminal services session bookmarks and can configure the system to create terminal services resource policies that enable access to the servers specified in the terminal services session bookmarks.

To specify general Terminal Services options:

1. In the admin console, choose **Users > User Roles > Role > Terminal Services > Options**.
2. If you are enabling Citrix sessions, under Citrix client delivery method, specify where the system should obtain the ICA client to download to users' systems:
 - Download from the Citrix website-The system installs the latest version of the ICA client from the Citrix web site. You can edit the URL to point to a new location if the one listed is no longer valid.
 - **Download from the IVE** - Use the Browse button to browse to the ICA client on your local network. You can upload a CAB, MSI or EXE file. Once you upload the client, the system uses it as the default and downloads it to your users' systems when necessary. You must also specify the exact version number of the ICA client.

If you upload an MSI or EXE file, an open/save dialog box appears to download and install the client. If Java fallback is configured, you are given the option to bypass this download and use Java instead.

- **Download from a URL** - The system installs the ICA client of your choice from the specified web site. You must also specify the exact version number of the ICA client. If Java fallback is configured, you are given the option to bypass this download and use Java instead.

NOTE:

We recommend that you test the Citrix client that you expect the system to download with the custom ICA file that you have uploaded to the system. Perform this testing without the system to determine if the Citrix client supports the features in the custom ICA file. If the features do not work without the system, they will not work through the system either.

If you choose to download an ICA client from the Citrix web site or a URL, the system secures the download transaction by processing the URL through the Content Intermediation Engine. Therefore, you must choose a site that is accessible by the system and enabled for users within this role.

To determine if the ICA web client is already installed on a machine, check for the following entry in your Windows registry: HKEY_CLASSES_ROOT\CLSID\{238F6F83-B8B4-11CF-8771-00A024541EE3}

You can determine the version number of an ICA client by extracting the cab file (for example, wficat.cab), looking for an inf file in the archive (for example, wficat.inf), and then locating the information about each ocx in the inf file. For example, wficat.inf (in wficat.cab) might contain the following information:

```
[wfica.ocx]
file-win32-x86=thiscab
clsid={238F6F83-B8B4-11CF-8771-00A024541EE3}
FileVersion=8,00,24737,0
[wfica32.exe]
file-win32-x86=thiscab
FileVersion=8,00,24737,0
```

In this case, "8,00,23737,0" is the file version. (Note that the version includes commas instead of periods.)

3. Enable the **User can add sessions** option to enable users to define their own terminal session bookmarks and to enable users to access terminal servers through the system browse bar on the home page. When you enable this option, the Add Terminal Services Session button appears on the Terminal Services page the next time a user refreshes the user console.
4. Select the **Deny single sign-on** for sessions added by user option if you do not want the user Add Terminal Service Session page to include the Authentication section used for single sign-on. This setting is disabled by default. When enabled, it disallows SSO for all user-added terminal services sessions, even if the user had previously configured SSO authentication credentials when that was permitted. This option adds a security measure to protect against exploitation of a security breach. If SSO is allowed, an attacker who gains access to a user's home page could gain access to the terminal services added by the user.
5. Enable the **Auto-allow role Terminal Services sessions** option to enable the system to automatically enable access to the resources defined in the terminal session bookmark (rather than having to create resource policies). Note that this only applies to role bookmarks, not user bookmarks.
6. You may not see the **Auto-allow** option if you are using a new installation or if an administrator hides the option.

If you want to allow users to enable access to local devices through the bookmarks they create, select from the following options in the Allow users to enable local resources defined below section:

- **Users can connect drives** - Enables the user to create bookmarks that connect the local drives to the terminal server, enabling the user to copy information from the terminal server to his local client directories.
- **User can connect printers** - Enables the user to create bookmarks that connect his local printers to the terminal server, enabling the user to print information from the terminal server to his local printer.

- **User can connect COM ports** - Enables the user to create bookmarks that connects his COM ports to the terminal server, allowing communication between the terminal server and the devices on his serial ports.
- **Allow Clipboard Sharing** - Enables the user to create bookmarks that shares the contents of the clipboard between the user's host computer and the terminal server. Due to the limitations in the pre-6.0 versions of the RDP client, disabling the Allow Clipboard Sharing option will automatically disable the Connect local drives, Connect local printers, and Connect COM Ports options.

When you enable local resources through the terminal server, each user can only access his own local resources. For instance, user 1 cannot see user 2's local directories.

The Connect Devices options that you specify at the role-level override any Connect Devices options that you set at the bookmark level.

- **User can connect smart cards** - Allows users to use smart card readers connected to their system for authenticating their remote desktop session.
- **User can connect sound devices** - Allows users to redirect audio from the remote desktop session to their local system.

NOTE:

- Smart cards redirecting audio are supported by Microsoft Remote Desktop Protocol versions 5.1 and later.
- If smart card option is selected, then Network Level Authentication (NLA) is not supported.
- **User can connect to Multiple Monitors** - Allows users to fully utilize all the monitors connected to the client computer for the remote desktop connection thereby providing extra desktop space and an almost seamless experience with the client desktop that is much improved over "Span mode".
- **User can connect microphone devices** - Allows users to use microphone devices connected to their system.

- **User can enable/disable NLA** - Allows users to enable/disable NLA at bookmark level.

Multiple monitors are supported by Microsoft Remote Desktop Protocol versions 8.1 and later.

7. In the Allow users to modify Display settings below section:

- Select **Desktop background** to display your current wallpaper setting. If you do not select this option, your background is blank.
- Select **Show contents of window while dragging** to show the contents of the Windows Explorer window while moving the window around your desktop.
- Select **Menu and window animation** to animate the movement of windows, menus, and lists.
- Select Themes to allow Windows themes to be set in the terminal server window.
- Select Bitmap Caching to improve performance by minimizing the amount of display information that must be passed over a connection.
- Select Font Smoothing (RDP 6.0 onwards) to make text smoother and easier to read. This option only works on Windows Vista computers running RDP clients that are version 6.0 or later.

8. Click Save Changes.

Configuring Terminal Services Resource Policies

When you enable the Terminal Services feature for a role, you need to create resource policies that specify which remote servers a user can access. You can create resource policies through the standard interface (as described in this section) or through resource profiles (recommended method).

The information in this section is provided for backwards compatibility. We recommend that you configure access to Windows terminal servers and Citrix servers through resource profiles instead, since they provide a simpler, more unified configuration method.

When writing a Terminal Services resource policy, you need to supply key information:

- Resources-A resource policy must specify one or more resources to which the policy applies. When writing a Terminal Services policy, you need to specify the terminal server to which users can connect.
- Roles-A resource policy must specify the roles to which it applies. When a user makes a request, the system determines what policies apply to the role and then evaluates those policies that correspond to the request.
- Actions-A Terminal Services resource policy either allows or denies access to a terminal server.

The system's engine that evaluates resource policies requires that the resources listed in a policy's Resources list follow a canonical format.

To write a Terminal Services resource policy:

1. In the admin console, choose Users > Resource Policies > Terminal Services > Access.
2. On the Terminal Services Policies page, click New Policy.
3. On the New Policy page, enter a name to label this policy and optionally description.
4. In the Resources section, specify the servers to which this policy applies.
5. In the Roles section, specify which roles to which this policy applies.
6. In the Action section, specify:
 - Allow access-To grant access to the servers specified in the Resources list.

- Deny access-To deny access to the servers specified in the Resources list.
 - Use Detailed Rules-To specify one or more detailed rules for this policy.
7. Click Save Changes.
 8. On the Terminal Services Policies page, order the policies according to how you want the system to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

[About Terminal Services Resource Profiles](#)

[Specifying Resources for a Resource Policy](#)

[Writing a Detailed Rule for Resource Policies](#)

Specifying the Terminal Services Resource Option

Use the Options tab to match IP addresses to hostnames specified as resources in your terminal services resource policies. When you enable this option, the system looks up IP address corresponding to each hostname specified in a Terminal Services resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

When you enable this option, the system compiles a list of hostnames specified in the Resources field of each Terminal Services resource policy. The system then applies the option to this comprehensive list of hostnames.

This option does not apply to hostnames that include wildcards and parameters.

To specify the Terminal Services resource option:

1. In the admin console, choose Users > Resource Policies > Terminal Services > Options.
2. Select IP based matching for Hostname based policy resources.
3. Click Save Changes.

Using the Remote Desktop Launcher

End users can connect to a terminal server by:

- Entering rdp://hostname in the system browser bar
- Creating a terminal services bookmark
- Using the remote desktop launcher (RDPLauncher)

RDPLauncher uses the Terminal Services section in the end-user home page and allows the end user to enter a terminal service IP address or hostname. The default server port is 3389.

RDPLauncher provides only the screen. User experience options are not available through RDPLauncher. For example, the following options in the New Terminal Services Sessions window do not apply to terminal services launched through RDPLauncher:

- Client port
- Authentication settings
- Start application settings
- Connect Devices settings
- Display Settings
- Remote Audio

To allow end users to use RDPLauncher,

1. Select the Terminal Services option in Users > User Roles > Role Name > General > Overview.
2. Select Enable Remote Desktop Launcher in Users > User Roles > Role Name > Terminal Services > Options.

3. (optional) If your end users are on non-Windows systems, such as a Macintosh or Linux system, select Enable Java for Remote Desktop Launcher and select the applet to use.



If you select Hob- RDP Applet from the Applet to Use menu, you must select the Configure HTML for the default applet check box in order to edit the HTML. Otherwise, the default HTML is used.

Screen size and color depth parameters for the RDPLauncher terminal services session are defined through Preferences > General on the end-users home page.

Remote Desktop and Telnet/SSH via HTML5 Access

HTML5 Access is a client-less solution to access Remote Desktops using Remote Desktop Protocol (RDP), or to connect to internal server hosts using Telnet protocols, or to communicate over an encrypted Secure Shell (SSH) session.

From 9.1R11 release onwards, Advanced HTML5 Access solution is released for General Availability (GA).

From 9.1R9 release, ICS has additional support for Advanced HTML5 Access solution on trial basis. This Advanced HTML5 Access solution supports two users by default. For additional users license, please contact Ivanti Support.



Advanced HTML5 Access solution is disabled when FIPS mode is turned ON and is enabled when FIPS mode is turned OFF. FIPS mode is applicable for the entire cluster.

Configuring the HTML5 Access Feature

Creating a HTML5 Access Resource Profile

A HTML5 Access resource profile is a profile that enables users to connect to Remote Desktops or to connect to internal server hosts in the clear using Telnet protocols or to communicate over an encrypted Secure Shell (SSH) session through a Web-based terminal session emulation.

Prior to 9.1R13, the HTML5 license count was based on the number of active HTML5 sessions. From 9.1R14, the HTML5 licenses will be counted based on the number of active users. Each user will be allowed to access up to five user sessions.

To create a HTML5 Access resource profile:

The following figure depicts Creating a HTML5 Access Resource Profile (**Users > HTML5 Access Resource Profiles > New HTML5 Access Resource Profile**):

HTML5 Access Resource Profiles >
New HTML5 Access Resource Profile

Solution Type: Basic HTML5 Advanced HTML5

Type: Windows RDP

Name: *

Description:

Host: * Name or IP address of host

Server Port: 3389

Create an access control policy for HTML5 access.

[Save and Continue >](#)

*indicates required field

1. In the admin console, choose **Users > Resource Profiles > HTML5 Access**.
2. Click **New Profile**.
3. Select the **Solution Type** as **Basic HTML5** or **Advanced HTML5**.
4. From the **Type** list, specify the session type (Windows RDP or SSH or Telnet) for this resource profile. If you have selected Advanced HTML5 solution type, you can also specify VNC session type.
5. Enter a unique name and optionally a description for the resource profile. (This name becomes the default bookmark's name.)
6. In the **Host** field, enter the Hostname, IP or user attribute of the server to which this resource profile should connect.
7. In the **Server Port** field, enter the port on which the system should connect to the server. (By default, the system populates this field with port number **3389** if you select Windows RDP, port number **23** if you select Telnet, port number **22** if you select SSH and port number **5900** if you select VNC.)
8. Select the **Create an access control policy for HTML5 Access** check box to enable access to the server specified in the Server Port box (enabled by default).
9. Click **Save and Continue**.

10. In the **Roles** tab, select the roles to which the resource profile applies and click **Add**.
The selected roles inherit the autopolicy and bookmarks created by the resource profile. If it is not already enabled, the system also automatically enables the HTML5 Access option in the Users > User Roles > Select Role > General > Overview page of the admin console for all of the roles you select.
11. Click **Save Changes**.
12. (Optional) In the **Bookmarks** tab, modify the default bookmark created by the system and/or create new ones. (By default, the system creates a bookmark to the server defined in the **Host** field and displays it to all users assigned to the role specified in the **Roles** tab.)

Defining Bookmarks for HTML5 Access Resource Profile

When you create a HTML5 Access resource profile, the system automatically creates a bookmark that links to the host that you specified in the resource profile. The system enables you to modify this bookmark as well as create additional bookmarks to the same host.

From 9.1R12, FIPS enabled users can create admin/end-user Advanced HTML5 bookmarks. However, Advanced HTML5 feature is still not FIPS compliant.

To define bookmarks for HTML5 Access resource profile:

1. In the admin console, select **Users > Resource Profiles > HTML5 Access > Resource Profile Name > Bookmarks**.
2. Click the appropriate link in the Bookmark column if you want to modify an existing session bookmark. Or, click **New Bookmark** to create an additional session bookmark.
Although it is generally easy to create a resource profile session bookmark through the resource profile configuration page, you can choose to create one through the user roles page as well.
3. (Optional.) Change the name and description of the session bookmark. (By default, the system populates and names the session bookmark using the resource profile name.)

The following figure depicts Creating a HTML5 Access Resource Profile - Bookmarks Configuration (**Users > Resource Profiles > HTML5 Access > Bookmarks**):

ivanti System Authentication Administrators **Users** Maintenance Wizards Ivanti Connect Secure

RDP
HTML5 Access Resource Profiles > Bookmarks >

Solution Type: Advanced HTML5
Connection Type: Windows RDP

Name: RDP

Description:

Host: * 10.96.158.23 Name or IP address of remote host

Port: * 3389

Options

Bookmark opens new window...

Authentication - Single Sign On

Username: administrator **Fetch Domain** Username or <USER> or Domain\Username

Variable Password: <PASSWORD> or <PASSWORD@SECAuthServer>

Password:

Screen Settings

Color Depth: 24 Number of bits to indicate color

Width: Desktop screen width: 800 min, 1920 max

Height: Desktop screen height: 600 min, 1080 max

DPI: Dots Per Inch

Resource Options

Disable Audio Enable remote drive for file transfer

Enable Printing Connect to the console session

Enable audio on console session Enable copy/paste

Enable Audio Recording Enable High Sound Quality

Enable Multiple Monitors Enable Session Recording

Enable Camera Store session record: End user External Storage

Performance Flags

Enable Wallpaper Enable Theming

Enable Font Smoothing Enable Full Window Drag

Enable Desktop Composition Enable Menu Animations

Other Settings

Keyboard Layout: English (US) keyboard

Encryption: Allow the server to choose

Remote Program Options

Program Type: Shell Program Remote App

Start program on connection:

Remote Dir:

Roles

Specify which user roles will get this bookmark.

ALL selected roles

Subset of selected roles...

Save changes?

Save Changes

*indicates required field

Licensed to 0312M154A0EP504VS
Copyright © 2022 Ivanti, Inc. All rights reserved.

- Allow users to open the bookmark in a new window by configuring the "Bookmark opens new window..." option and specifying how to display the browser address bar and browser toolbar.

5. Pass user credentials from the system to the terminal server so that users can sign onto the terminal server without having to manually enter their credentials. You can do this by configuring options in the Authentication - Single Sign On area of the bookmark configuration page.
6. Specify how the terminal emulation window should appear to the user during a terminal session by configuring options in the Screen Settings area of the bookmark configuration page.
7. Allow users to access local resources such as printers and drives through the terminal session by configuring options in the Resource Options area of the bookmark configuration page.
 - **Disable Audio** - To disable sound during the remote session.
 - **Enable Printing** - To grant access to the servers specified in the Resources list.
 - **Enable audio on console session** - To grant access to the servers specified in the Resources list.
 - **Enable copy/paste** - To grant copy/paste capability for particular resource.
 - **Enable remote drive for file transfer** - To grant access to the servers specified in the Resources list.
 - **Connect to the console session** - To grant access to the servers specified in the Resources list
 - ***Enable Audio Recording** - To grant access to the audio recording during the remote session.
 - ***Enable High Sound Quality** - To grant access to the high sound quality.
 - ***Enable Multiple Monitors** - To grant access for multiple monitors connected to the client computer during the remote session.
 - ***Enable Session Recording** - To grant access to the recording of end user sessions. Select the option to store the recordings to an End User system or to an external storage.
 - ***Enable Camera** - To grant access to the web camera.

* Options available for Advanced HTML5 solution.
8. Allow users to access specific applications on the terminal server by configuring options in the Remote Program Options area of the bookmark configuration page. In addition, you can use settings in this area to define auto-launch and application directory and arguments options.

9. Specify the roles to which you want to display the session bookmarks if you are configuring the session bookmark through the resource profile pages, under Roles:
 - **ALL selected roles** - Displays the session bookmark to all of the roles associated with the resource profile.
 - **Subset of selected roles** - Displays the session bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL Selected Roles list and click Add to move them to the Subset of selected roles list.
10. Click **Save Changes**.

When a user accesses a HTML5 RDP bookmark without SSO to access backend resources, the client prompts for credentials before opening the HTML5 session.



The client does not provide options to change password.

Creating HTML5 Enduser Bookmarks

The following figure depicts Creating a HTML5 Enduser Bookmark for Remote Desktop (**Users > Roles > Users**):



- With regard to an end user, if the **Allow user to add session** is enabled, an icon appears in the end user's page to add HTML5 access session. Options are similar to admin bookmark options based on the settings an admin allows a user to change.
- Options indicated with * are available for Advanced HTML5 bookmarks.

1. In the admin console, choose **Users > User Roles > Role > HTML5 Access > Options**.

The administrator has the option to select the solution type as basic or Advanced HTML for each user role. Basic HTML5 is selected by default.

2. Enable the "User can add sessions" option to enable users to define their own HTML5 Access session bookmarks. When this option is enabled, the Add HTML5 Access Session button appears on the HTML5 access panel the next time a user refreshes the user console. Admin can choose to enable RDP, SSH, Telnet or VNC sessions.

3. Enable Record sessions for bookmarks created by end users to record and store the session on an external storage.
4. Enable Remote Desktop Launcher to enable users to access HTML5 Access servers through the browse bar on the home page
5. Select the **Deny single sign-on for sessions added by user** option if you do not want the user Add HTML5 Access Session page to include the Authentication section used for single sign-on. This setting is disabled by default.
6. If you want to allow users to enable access to devices through the RDP bookmarks, select from the following options in the **RDP-Allow users to enable resources defined below** section:
 - **User can Disable Audio** - to disable sound during the remote session
 - **User can Enable remote drive for file transfer** - to connect the user's local drive to the terminal server, enabling the user to copy information from the terminal server to his local client directories
 - **User can Enable Printing** - to connect the user's local printers to the terminal server, enabling the user to print information from the terminal server to his local printer
 - **User can Connect to the console session** - to connect to the console (admin) session of the RDP server
 - **User can enable audio on console session** - to play the audio only at the server
 - **User can enable copy/paste** - to enable copy from the rdp session and paste to the clipboard
 - ***User can enable high sound quality** - to enable high sound quality
 - ***User can enable audio recording** - to enable audio recording of the user session
 - ***User can enable session recording** - to enable session recording of the user session. Select the
 - ***User can enable multiple monitor** - to enable maximum of four monitors connected to the client computer for the remote desktop connection thereby providing extra desktop space
 - ***User can enable camera redirection** - to enable web camera redirection

7. If you want to allow users to enable RDP performance flags through the bookmarks they create, select from the following options in RDP- Allow users to enable performance flags defined below section:
 - **User can enable wallpaper** - to allow users to display a wallpaper background to users
 - **User can enable theming** - to allow users to set Windows themes in their terminal server windows
 - **User can enable font smoothing** - to allow users to make text smoother and easier to read. This option only works on Windows Vista computers running RDP clients that are version 6.0 or later
 - **User can enable full window drag** - to enable users to specify the contents of the Internet Explorer window while they move the windows on their desktops
 - **User can enable desktop composition** - to allow users to make text smoother and easier to read. This option only works on Windows Vista computers running RDP clients that are version 6.0 or later
 - **User can enable menu animations** - to enable users to animate the movement of menus
8. If you want to allow users to enable access to devices through the SSH bookmarks, select from the following options in the **SSH-Allow users to enable resources defined below** section:
 - **User can enable SFTP** - to enable users to establish SFTP connections
 - ***User can enable Copy/Paste** - to enable copy from the session and paste to the clipboard
9. If you want to allow users to enable access to devices through the Telnet bookmarks, select from the following options in the **Telnet-Allow users to enable resources defined below** section:
 - ***User can enable Copy/Paste** - to enable copy from the rdp session and paste to the clipboard
10. If you want to allow users to enable access to devices through the VNC bookmarks, select from the following options in the **VNC-Allow users to enable resources defined below** section:
 - ***User can enable Copy/Paste** - to enable copy from the session and paste to the clipboard
 - **User can track remote cursor locally** - to enable rendering remote system cursor locally by the viewer

- **User can ignore remote cursor** - to enable ignoring the remote cursor

Configuring External Storage

The following figure depicts External Storage configuration.

The screenshot shows the Ivanti Connect Secure admin console. The top navigation bar includes 'System', 'Authentication', 'Administrators', 'Users', 'Maintenance', and 'Wizards'. The main content area is titled 'Storage Configuration' under 'HTML5 Profiles'. It features a checkbox for 'Enable external storage'. Below this are three input fields: 'Storage Path' (with a placeholder 'Samba share, smb://<name or ip-address>/<Share>'), 'Username' (with a placeholder 'Username or Domain\Username to access share'), and 'Password'. A 'Save Changes' button is located below the fields. A note at the bottom states: '*Note: This will disconnect all existing advanced HTML5 connections.'

To configure the external storage for session recordings:

1. In the admin console, navigate to **Users > Resource Profiles > HTML5 Access > Storage Configuration**.
2. Select **Enable external storage**.
3. Enter the complete storage path to store the session recordings.
4. Enter the Username and Password required to access the location.
5. Click **Save Changes**.

Defining SSO Options for the Remote Desktop Session

The following figure depicts Defining SSO Options for the Remote Desktop Session (**Users > Authentication - Single Sign On**):

The screenshot shows the 'Options' section of a Remote Desktop bookmark configuration page. At the top, there is a checkbox labeled 'Bookmark opens new window...'. Below this is the 'Authentication - Single Sign On' section. It contains three input fields: 'Username:' with a dropdown menu showing '<USER>', 'Variable Password:' with a dropdown menu showing '<PASSWORD>', and 'Password:' with an empty text box. To the right of the 'Variable Password' field is a blue button labeled 'Fetch Domain'. Below the 'Fetch Domain' button is the text 'Username or <USER> or Domain\Username'. Below the 'Password' field is the text '<PASSWORD> or <PASSWORD@SEcAuthServer>'. Below the authentication section is the 'Screen Settings' section. It contains three input fields: 'Color Depth:' with a dropdown menu showing '24', 'Width:' with an empty text box, and 'Height:' with an empty text box. To the right of the 'Color Depth' field is the text 'Number of bits to indicate color'. To the right of the 'Width' field is the text 'Desktop screen width: 800 min, 1920 max'. To the right of the 'Height' field is the text 'Desktop screen height: 600 min, 1080 max'. A red arrow points down to the 'Fetch Domain' button.

To define single sign-on options:

1. Create Remote Desktop bookmark or edit an existing bookmark.
2. Scroll to the Authentication - Single Sign On area of the bookmark configuration page.
3. Specify **Username** to pass to the terminal server. You can enter a static username or a variable. Enter the <username> variable to pass the username stored in the system's primary authentication server. Or use the following syntax to submit the username for the secondary authentication server: <username@SecondaryServerName> or <username[2]>. The Fetch domain is provided for the admins and end-user created bookmarks. This option helps to fetch the domain name from the remote AD machine.
4. Specify **Password** if you want to specify a static password or specify **Variable Password** if you want to use the password stored in the system's primary or secondary authentication server. To use the password from the primary authentication server, enter the <password> variable. Or use the following syntax to submit the password for the secondary authentication server: <Password@SecondaryServerName> or <Password[2]>.
5. Click **Save Changes**.

Defining Display Options for the Remote Desktop Session

When configuring Remote Desktop bookmark, you can specify how the terminal emulation window should appear to users during their terminal sessions.

To define display settings for the users' sessions:

1. Create a Remote Desktop bookmark or edit an existing bookmark.
2. Scroll to the Screen Settings area of the bookmark configuration page.

3. Select number of bits to indicate color in the **Color Depth** drop-down list. The default color depth is 24bit.
4. Enter the desktop screen width in the **Width** box. You can set it to minimum 800 and maximum 1920.
5. Enter the desktop screen height in the **Height** box. You can set it to minimum 600 and maximum 1080.
6. Enter the screen resolution in the **DPI** box.
7. Click **Save Changes**.

Defining Device Connections for the Remote Desktop Session

To define local resources that users can access:

1. Create a Remote Desktop bookmark or edit an existing bookmark.
2. Scroll to the Resource Options area of the bookmark configuration page.
3. Select **Enable remote drive for file transfer** to connect the user's local drive to the terminal server, enabling the user to copy information from the terminal server to his local client directories.
4. Select **Enable Printing** to connect the user's local printers to the terminal server, enabling the user to print information from the terminal server to his local printer.
5. Select **Disable Audio** to disable sound during the remote session. Select Enable audio on console session to play the audio only at the server.



- Sound options are supported by Microsoft Remote Desktop Protocol versions 5.1 and later.
 - File transfer (using the new HTML5/RDP feature) does not work if the Disable Audio option is checked.
-

6. If you want to allow users to enable performance flags through the bookmarks they create, select from the following options in Allow users to enable performance flags defined below section:

- **User can enable wallpaper** - to allow users to display a wallpaper background to users.

- **User can enable theming** - to allow users to set Windows themes in their terminal server windows.
 - **User can enable font smoothing** - to allow users to make text smoother and easier to read. This option only works on Windows Vista computers running RDP clients that are version 6.0 or later.
 - **User can enable full window drag** - to enable users to specify the contents of the Internet Explorer window while they move the windows on their desktops.
 - **User can enable desktop composition** - to allow users to make text smoother and easier to read. This option only works on Windows Vista computers running RDP clients that are version 6.0 or later.
 - **User can enable menu animations** - to enable users to animate the movement of menus.
7. Click **Save Changes**.
 8. For a detailed file transfer procedure, refer to the KB article: [File Transfer on Remote Desktop via HTML5 Access](#).

Defining Application Settings for the Remote Desktop Session

When configuring Remote Desktop bookmark, you can specify that users can only access specific applications on the terminal server.

To define applications that users can access:

1. Create Remote Desktop bookmark or edit an existing bookmark.
2. Scroll to the Remote Program Options area of the bookmark configuration page.
3. Specify the program that you want to launch automatically on connection in the **Specify program on connection box**.
4. Enter the application name (applicable only for servers running Windows 2008 and later) in the **Remote App** box.
5. Specify where the application's executable file resides on the terminal server in the **Remote App Dir** box (visible only when you clear Launch seamless window). For example, you might enter the following directory for the Microsoft Word application: C:\Program Files\Microsoft Office\Office10\WinWord.exe
6. Specify the arguments for the application in the **Remote App Args** box.



You can use session variables such as <username> and <password> in the Remote App Args box. For example, when specifying an application path, you might want to include the <username> variable to personalize the location. For example: C:\Documents and Settings\<>username>\My Documents.

7. Click **Save Changes**.



Windows requires a special notation for the names of remote applications. The names of remote applications must be prefixed with two vertical bars. For example, if you have created a remote application on your server for notepad.exe and have assigned it the name "notepad", you would set this parameter to: "||notepad".

Defining VNC Bookmarks for HTML5 Access Resource Profile

When you create a HTML5 Access resource profile with VNC session type, the system automatically creates a bookmark that links to the host that you specified in the resource profile. The system enables you to modify this bookmark as well as create additional bookmarks to the same host.

To define bookmarks for HTML5 Access resource profile:

1. In the admin console, select **Users > Resource Profiles > HTML5 Access > Resource Profile Name > Bookmarks**.
2. Click the appropriate link in the Bookmark column if you want to modify an existing session bookmark. Or, click **New Bookmark** to create an additional session bookmark.
3. (Optional.) Change the name and description of the session bookmark. (By default, the system populates and names the session bookmark using the resource profile name.)

The following figure depicts Creating an HTML5 Access Resource Profile - Bookmarks Configuration:

HTML5 Access Resource Profiles > Bookmarks >

VNC

Solution Type: Advanced HTML5
 Connection Type: VNC

Name:

Description:

Host: * Name or IP address of remote host

Port: *

Options

Bookmark opens new window...

Authentication - Single Sign On

Username: Username or <USER> or Domain\Username

Variable Password: <PASSWORD> or <PASSWORD@SECAuthServer>

Password:

VNC Settings

Color Depth: Number of bits to indicate color

Enable copy/paste
 Track remote cursor locally
 Ignore remote cursor

Encoding:

Roles

Specify which user roles will get this bookmark.

ALL selected roles
 Subset of selected roles...

Save changes?

*Indicates required field

4. Allow users to open the bookmark in a new window by configuring the "Bookmark opens new window..." option and specifying how to display the browser address bar and browser toolbar.
5. In the Authentication - Single Sign On section:
 - Specify **Username** to pass to the terminal server. You can enter a static username or a variable. Enter the <username> variable to pass the username stored in the system's primary authentication server. Or use the following syntax to submit the username for the secondary authentication server: <username@SecondaryServerName> or <username[2]>.
 - Specify **Password** if you want to specify a static password or specify **Variable Password** if you want to use the password stored in the system's primary or secondary authentication server. To use the password from the primary authentication server, enter the <password> variable. Or use the following syntax to submit the password for the secondary authentication server: <Password@SecondaryServerName> or <Password[2]>.
6. In the VNC Settings section:
 - Select number of bits to indicate color in the **Color Depth** drop-down list.

- Select **Enable Copy/Paste** option to grant copy/paste capability for particular resource.
 - Select **Track remote cursor locally** option to render remote system cursor locally by the viewer.
 - From the **Encoding** drop-down list, select the appropriate method for encoding the remote screen image.
7. Specify the roles to which you want to display the session bookmarks if you are configuring the session bookmark through the resource profile pages, under Roles:
- **ALL selected roles** - Displays the session bookmark to all of the roles associated with the resource- profile.
 - **Subset of selected roles** - Displays the session bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL Selected Roles list and click Add to move them to the Subset of selected roles list.
8. Click **Save Changes**.

When a user accesses a HTML5 VNC bookmark without SSO to access backend resources, the client prompts for credentials before opening the HTML5 session.

Remote Desktop User Experience

When you enable the Remote Desktops via HTML5 Access for a user role, the end user needs to specify the resource that the user wants to access and enter credentials for the resource.

Users can access remote desktop resources using the following methods:

- **URLs from other web sites** - In most cases, users access session bookmarks directly from the end-user console. If you do not want to require users to sign into the end-user console to find and access Remote Desktop links, you can create URLs on other web sites that point to session bookmarks that you have already created.
- **Connect Secure browse bar** - In addition to enabling users to link to Remote Desktop links through bookmarks and URLs, you can also enable them to access these resources through the system browse bar on Windows systems. Users can access Microsoft terminal services or remote desktop sessions by entering `hrdp://hostname` in the browse box.
- **Server address** - By entering the Remote Desktop IP address or hostname, users can launch a remote desktop connection to any accessible server.

Telnet/SSH User Experience

The HTML5 Access feature supports the following applications and protocols:

- **Network Protocols** - Supported network protocols include Telnet and SSH.
- **Terminal Settings** - Supported terminal settings include VT100, VT320, and derivatives and screen buffers.
- **Security** - Supported security mechanisms include Web/client security using SSL and host security (such as SSH if desired).

You can create secure terminal session bookmarks that appear on the welcome page for users mapped to a specific role. A terminal session bookmark defines Terminal Session information for Telnet or SSH sessions that users may launch. These sessions give users access to a variety of networked devices, including UNIX servers, networking devices, and other legacy applications, that utilize terminal sessions. The system supports SSH versions V1 and V2 and uses the following SSH versions: OpenSSH 5.2, OpenSSH_2.9.9p1, SSH protocols 1.5/2.0, and OpenSSL 0x0090607f.

Monitoring HTML5 Sessions

In 9.1R7, the current HTML5 sessions information is provided in Dashboard and the trend graph. This information helps administrator to view the CPU usage and take necessary action to provide better remote access experience for the users. The connection type is logged as HTML5.

To enable HTML5 graph:

1. Select **System > Status > Overview**.
2. In the **Select list of graphs** list, enable the **HTML5 Connections** option. By default, this option is enabled.

The HTML5 Connections graph shows the traffic on the HTML5 RDP, HTML5 SSH, and HTML5 Telnet connections.

3. Select **System > Status > Virtual Desktop Sessions**.

The Active Virtual Desktops Sessions page lists the active user sessions and the connection types.

4. Select **System > Log Monitoring > User Access > Log** to view the HTML5 sessions log.

Launching Custom Page via HTML5 Access

An end user can launch either Basic HTML5 session or Advanced HTML5 session. End users can connect to a target server by entering the following in the browser bar:

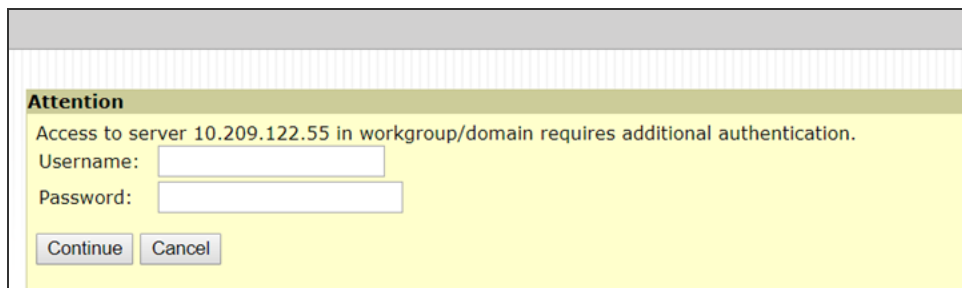
```
https://<PCS-FQDN>/dana/html5acc/html5urllaunch.cgi?type=launcher&host=<TargetMachineIP>
&port=3389&stype=0&width=600&height=480&dpi=96&security=tls&enable-
wallpaper=true&enable-full-windowdrag=true&
username=admin&password=pcs123&enable-drive=false&enable-
printing=true&disable-audio=true&client-name=<any-string>
```

To allow end users to use RDPLauncher:

1. Navigate to **Users > User Roles > Role Name > General > Overview** and select the **HTML5 Access** option.
2. **Navigate to Users > User Roles > Role Name > HTML5 Access > Options and do the following:**
 1. Select **Enable Remote Desktop Launcher**.
 2. Select necessary resources which user wants to access.
 3. Select necessary performance flags which user wants to access.

If the user is not logged in to ICS, it will prompt for ICS login and then prompt for target server credentials as shown in the screenshot below. Upon providing necessary details, it will open the HTML5 session.

The following figure depicts Additional Authentication in the Target Server:



The parameter can be validated from the RDP client task manager -> Users > client name.

Parameters that can be configured via query parameters are:

- disable-audio (true/false)

- enable-drive (true/false)
- enable-printing (true/false)
- console (true/false)
- console-audio (true/false)
- enable-wallpaper (true/false)
- enable-theming (true/false)
- enable-font-smoothing (true/false)
- enable-full-window-drag (true/false)
- enable-desktop-composition (true/false)
- enable-menu-animations (true/false)
- color-depth(8/16/24)
- security (rdp, nla, tls and any)
- server-layout(en-us-qwerty, de-de-qwertz,fr-fr-azerty, it-it-qwerty, sv-se-qwerty, failsafe)
- color-scheme (black-white, white-black, gray-black, green-black)
- font-name (courier, monospace etc...)
- font-size
- width
- height
- dpi
- host
- port
- stype (eg: 0=rdp, 1=ssh and 2 = telnet)
- ignore-cert (true)
- client-name

VPN Tunneling

The VPN tunneling access option (formerly called Network Connect) provides a VPN user experience, serving as an additional remote access mechanism to corporate resources using Connect Secure. This feature supports all Internet-access modes, including dial-up, broadband, and LAN scenarios, from the client machine and works through client-side proxies and firewalls that allow SSL traffic.

When a user launches VPN tunneling, the system transmits all traffic to and from the client over the secure VPN tunnel. The only exception is for traffic initiated by other system-enabled features, such as Web browsing, file browsing. If you do not want to enable other system features for certain users, create a user role for which only the VPN tunneling option is enabled and make sure that users mapped to this role are not also mapped to other roles that enable other system features.

For details about the configuration, refer to the *Ivanti VPN Tunneling Configuration Guide* available on the <https://www.ivanti.com/support/product-documentation> site.

Enterprise Onboarding

Configuring Enterprise Onboarding

Enterprise onboarding allows users to securely access enterprise network resources with almost any device. Wi-Fi, VPN and certificate profiles can be defined for enterprise resources and downloaded to a device during onboarding, depending on the device type.

The profiles can be defined on a single Ivanti Connect Secure or Ivanti Policy Secure server dedicated to onboarding or they can be defined on each server. Alternatively, the profiles can be defined on a third-party MDM server, in which case users will see a link and instructions on the onboarding page to continue onboarding using the external MDM server.

Onboarding is initiated from the browser. The supported profiles depend on the device type and whether the Pulse client is installed (see the following table).

The following table lists the Device Onboarding Profile Support:

Device	Supported profiles
Android 4.0 or later	Supports all profiles, but the Pulse client must be installed during onboarding.
iOS 6.0 or later	Supports all profiles (Safari browser).
Windows 7.0, 8.0, and 8.1	Supports Wi-Fi and certificate profiles (IE, Firefox, or Chrome browser). The Pulse client onboarding application must be installed during onboarding. Windows 8 RT and Windows 8 Phone are not supported.
MAC OS X	Supports Wi-Fi and certificate profiles (Safari browser).

Enterprise onboarding is enabled in the user role, and each profile can be applied to all user roles or specific roles. The SCEP server and CSR templates allow certificates to be generated dynamically for device and server authentication.

- [Enabling Enterprise Onboarding at the Role Level](#)
- [Defining the SCEP Server](#)
- [Defining VPN Profiles](#)
- [Defining VPN Profiles](#)
- [Defining Wi-Fi Profiles](#)

- [Defining Certificate Profiles](#)
- [Onboarding Devices](#)
- [Workflow for Onboarding Android Devices](#)

Domain Discovery Service

In the email-based authentication, once user enters email, client would parse domain and send it to a discovery server to fetch the server URL. It would then proceed with AD authentication with the server.



To set up the Auto-Discovery experience, you will need to contact Technical support through DevOps ticket. Once the needed information is provided (and validated), Technical Support will enable the Auto-Discovery experience for your Email Domain.

Enabling Enterprise Onboarding at the Role Level

To enable enterprise onboarding for a user role:

1. In the admin console, choose **Users > User Roles > RoleName > General > Overview**.
2. In the Enterprise Device Onboarding section, select the **Enterprise Onboarding** check box.
3. Click **Save Changes**.
4. Click the **Enterprise Onboarding** tab or click **Options** next to the **Enterprise Onboarding** check box to specify the following:
 - **Auto launch** - Displays the onboarding page when the user logs in to Ivanti Connect Secure or Ivanti Policy Secure if enterprise onboarding is enabled for the user's role (the default). If this option is disabled, an onboarding link is displayed on the home page.
 - **Use third party MDM for onboarding** - Displays a link on the onboarding or home page where the user can download profiles from an MDM server. Enter the URL for the MDM page in the text box.
 - **Install Pulse Client**: Enabling this option will automatically install Pulse client during onboarding from Windows OS
5. Click **Save Changes**.

Defining the SCEP Server

The Simple Certificate Enrollment Protocol (SCEP) server configuration and CSR templates allows each client device to dynamically obtain certificates for authentication.

To define the SCEP server:

1. In the admin console, choose **Users > Enterprise Onboarding**.
2. Specify the following information:

Setting	Description
SCEP Server URL	Enter the URL for a SCEP server. The following SCEP servers are supported: Microsoft AD 2008 Symantec mPKI
Challenge	Specify the password required by the SCEP server.
Retries	Specify the number of attempts to access the server when the first attempt fails.
Retry Delay	Specify the number of seconds between retry attempts.
Upload Encryption Certificate	Click Browse to upload the certificate used to encrypt SCEP requests. To upload the certificate automatically, select the Test Enrollment check box, select a CSR template, and click Test Configuration. To create a CSR template, see Defining CSR Templates

3. Click **Save Changes**.

Defining CSR Templates

If the SCEP server is configured, the Certificate Signing Request (CSR) templates can be used in the VPN, Wi-Fi, and certificate profiles to allow each onboarded device to dynamically obtain certificates for authentication on all mobile devices. Up to 10 templates can be defined.




All LDAP attributes (such as <ldap.userAttrName>) and variables (such as <user>) can be used in the Subject DN, Email, and Subject Alternative Name Value fields. However, if you enter an LDAP variable with a string vector data type in the Subject Alternative Name Value field, only the first value in the string will be used.

To define CSR templates:

1. In the admin console, choose **Users > Enterprise Onboarding > CSR Templates**.

2. To add a template, click **New CSR Template** or select an existing template that you want to change, duplicate, or delete. Clicking **Duplicate** creates a copy of the selected template with **Copy** of before the template name.
3. Specify the following information:

Setting	Description
Name	Specify the template name displayed in the list of CSR templates.
Subject DN	Specify the subject distinguished name. For example: CN= <USERNAME>,OU=Engineering=Ivanti All LDAP attributes (such as <ldap.userAttrName>) and variables (such as <user>) can be used.
Email	(Optional) Specify an email address with the <USER> variable, such as <USER>@ivanti.com.
Subject Alternative Name Type	Select an alternative name type if the CA requires an alternative subject name. The types include RFC-822 Name (an e-mail address), DNS domain name, URI, and IP address.
Subject Alternative Name Value	Specify one or more values for the selected alternative name type. Multiple values must be separated by a comma or space. <hr/>  If an LDAP variable is specified that has a string vector data type, only the first value in the string will be used. <hr/>
Key Size	Select the key size used by the SCEP server.

4. Click **Save Changes**.



1. The number of keys available in the system can be viewed at Users->Enterprise OnBoarding->CSR Templates.
2. The keys are generated only if:
 - A) The onboarding license is installed
 - B) A CSR template is configured for that key size
3. The max number of keys of each type is minimum of 10K keys and the number of onboard user license installed.
4. Key generation is CPU intensive and time consuming. If bulk users are going to onboard it is recommended to make sure that the number of available keys \geq the numbers of users that needs to onboard.

Defining VPN Profiles

VPN profiles provide Android and iOS devices with secure access to enterprise networks. One or more VPN profiles can be assigned to specific user roles or to all roles. Up to 10 profiles can be defined.



All LDAP attributes (such as <ldap.userAttrName>) and variables (such as <user>) can be used in the Username, Realm, and Role fields.

To define VPN profiles:

1. In the admin console, choose Users > Enterprise Onboarding > VPN Profiles.
2. To add a profile, click New Profile or select an existing profile that you want to change, duplicate, or delete. Clicking Duplicate creates a copy of the selected profile with Copy of before the profile name.
3. Specify the following profile information:

Setting	Description
Name	Specify the name to be displayed in the list of VPN profiles.
Description	(Optional) Enter a description of the VPN profile.
Apply to Client Types	Select the device types the profile applies to (Android and iOS only).
Server URL	Specify the URL of the VPN server (must be a Ivanti Connect Secure or Ivanti Policy Secure device).

Setting	Description
Realm	Specify the realm name. The realm is required only if the sign-in URL has the User picks from a list of authentication realms option enabled.
Role	Specify the user role. The user role is required if the role mapping rules for the user realm specify multiple roles and the User must select from among assigned roles option is enabled.
Username	Specify the <USER> variable for the user name.
Authentication Method	<p>Select Password or Certificate for the user authentication method. For certificate authentication, specify the following:</p> <p>Use CSR Template-Select the CSR template used to obtain the certificate. To create a CSR template, see "Defining CSR Templates".</p> <p>Enable VPN On Demand-Select this option to allow iOS devices to establish the VPN when a specific host or domain is accessed. To specify the first host or domain:</p> <p>Match Domain or Host-Enter a hostname or a partial domain name. For example, if you enter example.com, a match occurs when the user accesses any domain that ends with example.com, such as www.test-example.com.</p> <p>On Demand Action-When a match occurs on the specified host or domain, select whether a VPN is always established, never established, or only if the DNS look-up fails (Establish If Needed). Selecting Never Establish does not prevent an existing VPN from being used.</p> <p>To add another domain, click the + button. To remove a domain, select the check box next to the domain and click the - button. Up to 10 domains can be defined.</p>
Roles	<p>Select one of the following options:</p> <p>Policy applies to ALL roles-To apply this profile to all users.</p> <p>Policy applies to SELECTED roles-To apply this profile only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.</p> <p>Policy applies to all roles OTHER THAN those selected below-To apply this profile to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.</p>

4. Click **Save Changes**.

Defining Wi-Fi Profiles

Wi-Fi profiles provide Android, iOS, MAC OS X, and Windows devices with secure access to wireless networks. One or more Wi-Fi profiles can be assigned to specific user roles or to all roles. Up to 10 profiles can be defined.



All LDAP attributes (such as **<ldap.userAttrName>**) and variables (such as **<user>**) can be used in the Username and Password fields for the WPA Enterprise and WPA2 Enterprise security types.

To define Wi-Fi profiles:

1. In the admin console, choose **Users > Enterprise Onboarding > WiFi Profiles**.
2. To add a profile, click New Profile or select an existing profile that you want to change, duplicate, or delete. Clicking **Duplicate** creates a copy of the selected profile with **Copy** of before the profile name.
3. Specify the following profile information:

Setting	Description
Name	Specify the name to be displayed in the list of Wi-Fi profiles.
Description	(Optional) Enter a description of the profile.
Apply to Client Types	Select the device types the profile applies to (Android, iOS, MAC OS X, and Windows).
SSID	Specify the server set ID of the wireless network.
Non-Broadcast SSID	Select the check box if the wireless network does not broadcast its identity.
Auto Connect	Select the check box to connect the client automatically when the network is detected (not supported by Android clients).

Setting	Description
Security Type	<p>Select the type of authentication used by the network, and specify the password or enterprise settings, as required:</p> <p>None-No authentication required.</p> <p>WEP-Wired Equivalent Privacy used for a non-enterprise network. Enter the network shared key in the displayed text box.</p> <p>WPA Personal or WPA2 Personal-Wi-Fi Protected Access used for a non-enterprise network. You can select the encryption method (AES or TKIP) and enter the network shared key in the displayed text box (applies to Windows clients only).</p> <p>WPA Enterprise or WPA2 Enterprise-Wi-Fi Protected Access used for an enterprise network. Select the Extensible Authentication Protocols (EAP) supported by the network's RADIUS authentication server.</p> <p>For Android devices, note the following: Android 4.3 or later is required For the EAP-TLS protocol, the CA certificate must be configured (along with the client certificate) on Samsung devices for authentication. An 802.1x RADIUS server certificate must be signed by a private root CA. Authentication fails if the certificate is signed by an intermediate root CA.</p>
EAP	<p>For the WPA Enterprise and WPA2 Enterprise security types, select the supported EAP protocols and specify the associated authentication settings:</p> <p>None-If none of the EAP protocols is selected (Android devices only), enter the <USER> and <PASSWORD> variables in the Username and Password fields.</p> <p>iOS, MAC OS X, and Windows clients require at least one of the EAP types to be selected (PEAP, EAP-TLS, or EAP-TTLS).</p> <p>Selecting Multiple EAP types is not supported for Android clients.</p>

Setting	Description
PEAP	<p>The PEAP protocol is supported by all clients. Specify the following:</p> <p>Inner Authentication Method-Select the protocol used to authenticate the username and password (None or MSCHAPv2). The None option is valid only for Android devices.</p> <p>Username and Password-Enter the <USER> and <PASSWORD> variables.</p> <p>Outer Identity-Specify an alternate username to be used outside the encrypted tunnel, to conceal the user's identity in unencrypted packets.</p> <p>Trusted Server Name(s)-Specify the IP address or fully qualified domain name of one or more trusted RADIUS servers used by the network. Multiple servers must be separated by a semicolon.</p> <p>Trusted CA Certificate-For Windows clients, select the Trusted Root CA of the RADIUS server certificate, even if the device certificate is signed by an intermediate CA. The Trusted Root CA must be configured in a certificate profile before it can be selected here (see "Defining Certificate Profiles". For iOS and MAC OS X clients, if the RADIUS server certificate is signed by an intermediate CA, create a certificate profile for the intermediate CA, and then select the certificate here. The certificate profile ensures that the intermediate CA is downloaded to the client.</p>
EAP-TLS	<p>The EAP-TLS protocol is supported by all clients. Specify the following:</p> <p>Username-Enter the <USER> variable.</p> <p>Use CSR Template-Select the CSR template used to obtain the certificate. To create a CSR template, see Defining CSR Templates</p> <p>Trusted Server Name(s)-Specify the IP address or fully qualified domain name of one or more trusted RADIUS servers used by the network. Multiple servers must be separated by a semicolon.</p> <p>Trusted CA Certificate-For Windows clients, select the Trusted Root CA of the RADIUS server certificate, even if the device certificate is signed by an intermediate CA. The Trusted Root CA must be configured in a certificate profile before it can be selected here (see "Defining Certificate Profiles". On Windows 7 clients that have multiple certificates, users are prompted to select the certificate for 802.1x connections that use EAP-TLS.</p> <p>For iOS and MAC OS X clients, if the RADIUS server certificate is signed by an intermediate CA, create a certificate profile for the intermediate CA, and then select the certificate here. The certificate profile ensures that the intermediate CA is downloaded to the client.</p>

Setting	Description
EAP-TTLS	<p>The TTLS protocol is supported by all clients. Specify the following:</p> <p>Inner Authentication Method-Select the protocol used to authenticate the username and password (None, PAP, or MSCHAPv2). The None option is valid only for Android devices.</p> <p>Username and Password-Enter the <USER> and <PASSWORD> variables.</p> <p>Outer Identity-Specify an alternate username to be used outside the encrypted tunnel, to conceal the user's identity in unencrypted packets.</p> <p>Trusted Server Name(s)-Specify the IP address or fully qualified domain name of one or more trusted RADIUS servers used by the network. Multiple servers must be separated by a semicolon.</p> <p>Trusted CA Certificate-For Windows clients, select the Trusted Root CA of the RADIUS server certificate, even if the device certificate is signed by an intermediate CA. The Trusted Root CA must be configured in a certificate profile before it can be selected here (see Defining Certificate Profiles). Also, if the RADIUS server certificate is signed by an intermediate CA, then the public intermediate CA must be configured in a certificate profile to ensure that the intermediate CA is downloaded to the client along with the Wi-Fi TTLS profile configuration.</p> <p>For iOS and MAC OS X clients, if the RADIUS server certificate is signed by an intermediate CA, create a certificate profile for the intermediate CA, and then select the certificate here. The certificate profile ensures that the intermediate CA is downloaded to the client.</p> <p>Profile generation does not occur when Wi-Fi profile with EAP-TTLS is selected for windows 7 client. However, this issue is not seen with windows 8.1.</p>
Roles	<p>Select one of the following options:</p> <p>Policy applies to ALL roles-To apply this profile to all users.</p> <p>Policy applies to SELECTED roles-To apply this profile only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.</p> <p>Policy applies to all roles OTHER THAN those selected below-To apply this profile to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.</p>

4. Click **Save Changes**.

Defining Certificate Profiles

Certificate profiles specify the device certificates sent to each client device during onboarding. Up to 10 profiles can be defined.



For security reasons, certificate profiles cannot be included in the XML export or import.

To define certificate profiles:

1. In the admin console, choose **Users > Enterprise Onboarding > Certificate Profiles**.
2. To add a profile, click **New** Profile or select an existing profile that you want to change, duplicate, or delete. Clicking **Duplicate** creates a copy of the selected profile with Copy of before the profile name.
3. Specify the following information:

Setting	Description
Client Types	Select the device types the profile applies to (Android, iOS, MAC OS X, and Windows).
Import and Use Global Certificate	Select this option to use the Ivanti Connect Secure or Ivanti Policy Secure global certificate to authenticate the client device. Click Import Certificate & Key, click Browse to locate the certificate file, and then click Import. For more information about device certificates, see Using Device Certificates .
Import and Use CA Certificate	Select this option to import any CA certificate (public Root CA, private Root CA, public intermediate CA, or private intermediate CA). These CA's can be used in Wi-Fi profiles and must be downloaded to the client devices. Click Import and Use CA Certificate , click Browse to locate the certificate, and then click Import CA Certificate .
Generate per User Certificate	Select this option to use the SCEP server and a CSR template to generate a certificate for each client. Select a CSR template from the Use Certificate Template list. To create a CSR template, see Defining CSR Templates .

Setting	Description
Roles	<p>Select one of the following options:</p> <p>Policy applies to ALL roles-To apply this profile to all users.</p> <p>Policy applies to SELECTED roles-To apply this profile only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.</p> <p>Policy applies to all roles OTHER THAN those selected below-To apply this profile to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.</p>

4. Click **Save Changes**.

Onboarding Devices

Onboarding is initiated from the browser. When a user logs in, the onboarding option is displayed if VPN, Wi-Fi or certificate profiles are defined in the user's role. MAC OS X devices and iOS devices can be onboarded without installing the Pulse client. For Android devices the browser displays a link to install the Pulse client. For Windows devices, the browser displays a link to install the Onboarding application.

Note the following:

- If the device has a VPN connection to Ivanti Connect Secure, the user is warned that the connection will be closed and reestablished through the onboarding process.
- If the user onboards the device again, which may be necessary if a certificate expires or the configuration is deleted, new profiles are downloaded to the device.
- The following message IDs in the User Access Log can be used to verify the onboarding process (they include the username and device ID):
 - AUT31186-Indicates the status of an onboarding attempt (failed or successful)
 - AUT31152-Indicates onboarding failed because the maximum device limit of 10000 has been reached
 - AUT31187-Indicates the attempt to build a configuration profile failed due to an error
 - AUT31188-Indicates a configuration profile was generated successfully, and lists the names of the profiles contained in the configuration profile

The following message IDs are related to device limits:

- SYS31177-Indicates the number of devices onboarded is nearing the system limit of 10000.


- SYS31178-Indicates the number of devices onboarded has exceeded the system limit of 10000 (critical error).
- SYS31193-Information message generated by a background process that attempts to delete device records when 95% of system limit (10000) is reached. It displays the number of device records deleted, the current number of onboarded devices, and the system limit.

Workflow for Onboarding Android Devices

To onboard an Android device:

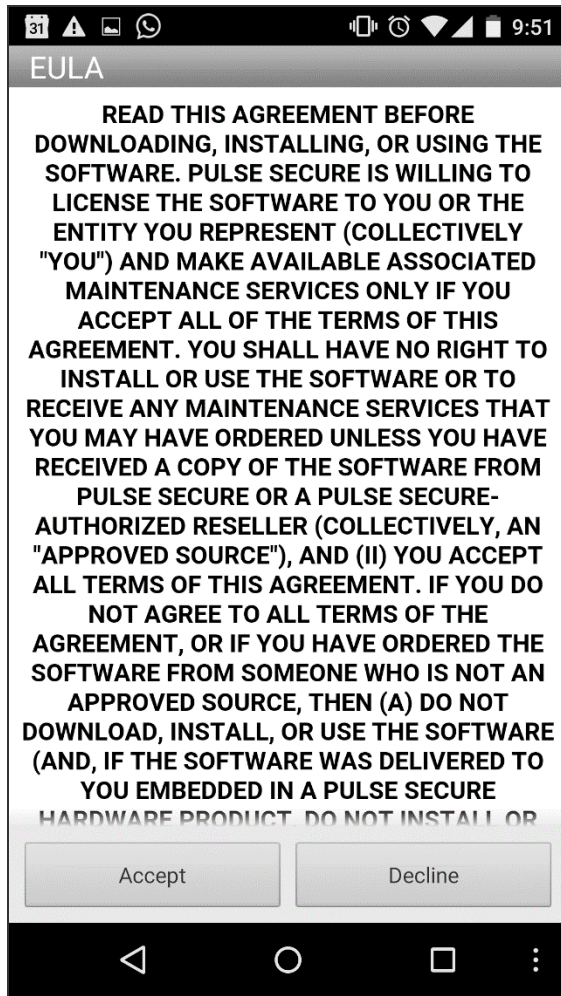
1. Enter the onboarding URL in the browser.

The following figure depicts the Onboarding Start Page:




2. The first time Pulse client is launched, the EULA is displayed.

The following figure depicts the End User License Agreement:




3. On the Pulse client onboarding Wizard page, read the instructions carefully and tap Next.

The following figure depicts the Onboarding Wizard Start Page:




4. On the CA certificate provisioning page, tap OK.

The following figure depicts the Certificate Provisioning Page:




5. Paste the password from the clipboard to extract the certificates, and tap OK.

The following figure depicts the Certificate Password Page:




6. Tap **OK** to confirm the certificate name.

The following figure depicts the Certificate Name Page:




7. View the number of certificates and profiles provisioned on the client, and tap Finish.

The following figure depicts the Onboarding Wizard Summary Page:



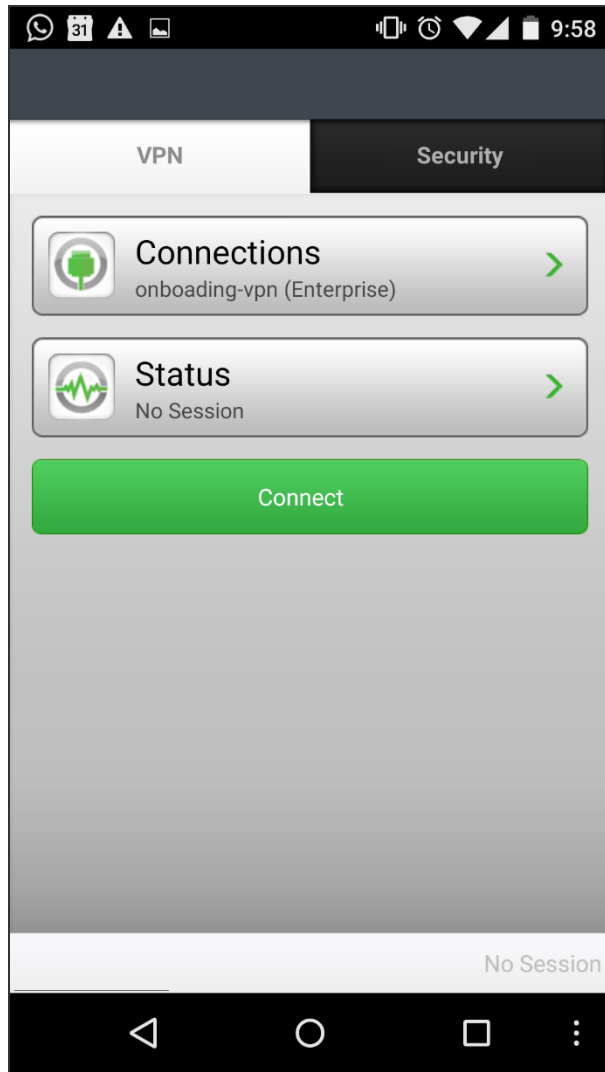
8. If onboarding fails due to an error, tap Retry. Users should contact their administrator if onboarding fails after three attempts.

The following figure depicts the Error Page:




9. If onboarding is successful, tap the VPN tab to view the provisioned VPN connections.

The following figure depicts the VPN Connections Page:



10. Tap the Wi-Fi icon to view the provisioned Wi-Fi networks. To enable a Wi-Fi connection, select the network and tap the Connect icon.

The following figure depicts the Wi-Fi Connections Page:



Managing Onboarded Devices

The Device Management page lists the following types of devices:

- **Onboarded devices**- Devices that have Enterprise Onboarding enabled in the user's role and have been onboarded during device registration. After a device is onboarded, it is displayed on the Device Management page until it is deleted.

The username, user roles, operating system, and registration date are shown for each device, along with the onboarded, and access status. Devices that become inactive or invalid must be deleted manually.

To view the Device Management page:

1. Select **System > Status > Devices**.
2. Use the controls described in Table 100 to view and manage the devices.

Cloud Secure

Cloud Secure provides secure, seamless, and compliant access to cloud resources on a hybrid IT environment where companies are combining the best of the cloud with their own localized data centers. Cloud Secure solution integrates with multiple Ivanti products such as Ivanti Connect Secure, Ivanti Policy Secure, Pulse Workspace etc.

Cloud Secure provides great level of flexibility with integration to various Third-Party vendors such as MDM vendors, IdP vendors etc.

It is a licensed feature, so the Administrator should procure and install the required license.

For details about the configuration, various deployment scenarios, reports, etc. refer to Cloud Secure documentation available at the <https://www.ivanti.com/support/product-documentation> site.

Network and Host Administration

Network and Host Administration Overview

When you install and initially set up the device, you use the serial port console to set basic network and host settings. To get started, you must use the serial console to configure these settings for the internal interface. You have the option to use the serial console to configure network and host settings for the external interface and the management interface. The network and host settings you configure with the serial port console include:

Once the internal interface has been configured, you can use the admin console Network Settings pages to modify settings for the internal interface, to enable and configure the external interface and the management interface, and to configure or manage advanced networking features, including:

- Hostname
- IPv6 addresses
- VLAN ports
- Virtual ports
- Route table entries
- Host mapping table entries
- ARP cache entries
- Neighbor discovery cache entries
- System date and time (manual configuration) or NTP

Configuring the Internal Port

The internal port connects to the local area network (LAN). The internal port settings are configured when you run the setup wizard from the serial console as part of the installation procedure. You can use the System > Network pages to make changes to the configuration.

To modify the internal port configuration:


1. Select **System > Network > Internal Port > Settings** to display the configuration page.

The following figure shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration as described in the following table.

3. Save your changes.

The following figure depicts the Ivanti Connect Secure Internal Port Configuration Page:



The following table lists the Internal Port Configuration Guidelines:

Settings	Guidelines
IPv4 Settings	
IP Address	Assign an IP address. You must assign an IPv4 address to the internal interface. An IP address is an identifier for a computer or device on a TCP/IP network. Networks using the TCP/IP protocol route messages based on the IP address of the destination. The format of an IPv4 address is a 32-bit numeric address written as four numbers separated by periods. Each number can be 0 to 255.

Settings	Guidelines
Netmask	Assign a netmask. A netmask indicates which part of an IP address indicates network identification and which part indicates the host identification. For example, the IP address and netmask 10.20.30.1 255.255.255.0 (or 10.20.30.1/24) refer to all the hosts in the 10.20.30.0 subnet. The IP address and netmask 10.20.30.1 255.255.255.255 (or 10.20.30.1/32) refer to a single host.
Default Gateway	Specify the IPv4 address for the default gateway for the routing domain to which the device belongs. A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.
IPv6 Settings	
Enable IPv6 / Disable IPv6	Disabled by default. Enable to support access from IPv6 endpoints. When you enable IPv6, the system acquires a link local address. If you switch from enabled to disabled, the system clears the link local address.
Link Local Address	Display the autoconfigured link local address (after you have enabled and saved the IPv6 configuration).
IPv6 Ad-dress	Specify a routable IPv6 address, such as a global unicast address that your network plan has provisioned for this host and interface. Automatic configuration methods are not supported. You must specify the appropriate address manually.
Prefix Length	Specify how many of the higher order contiguous bits of the IPv6 address comprise the prefix (the network portion of the IPv6 address). The default is 64.
Gateway	Specify the IPv6 address for the default gateway for the routing domain to which the device belongs. A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.
Advanced Settings	
MAC Ad-dress	Display the MAC address for the interface.
Link Speed	Specify the speed and duplex combination for the interface. If you run SNMP_GET and then change the Link Speed value, you must wait at least 5 minutes after submitting the change before running SNMP_GET again.

Settings	Guidelines
ARP Ping Timeout	<p>(IPv4 only.) Specify how long the system should wait for responses to Address Resolution Protocol (ARP) requests before timing out. Cluster nodes send ARP requests to the gateways of other nodes to determine if they are properly communicating with one another.</p> <p>If you have not deployed a cluster, the system does not use this setting. If the node belongs to a cluster, the timeout interval that you specify is synchronized across the cluster. In multisite clusters, you can override this setting for the individual nodes in the cluster using options in the System > Clustering page. Use caution when changing this setting in active/passive clusters, however, because the system also uses the ARP Ping Timeout setting on the Internal tab as a fail-over timer for the VIP.</p>
MTU	<p>Specify the maximum transmission unit.</p> <p>If IPv6 is enabled, the valid range is 1280 to 1500. If IPv6 is not enabled, the valid range is 576 to 1500.</p> <p>We recommend you retain the default MTU setting (1500) unless you must change the setting for troubleshooting purposes.</p> <p>If the administrator sets ignore-tcp-mss in Advanced Client Configuration, then the TCP MSS option is ignored during the virtual adapter MTU calculation on the client side. For details, see Using the Advanced Client Configuration Feature</p>
Default VLAN ID	<p>(Optional) Specify the default VLAN ID for the traffic of this port. When this parameter is set, all the traffic on this interface is subsequently tagged with the set VLAN ID and also accepts only incoming traffic with the same tag. Necessary changes are required on the connected switch port to handle bi-directional tagged traffic.</p> <p>If default VLAN ID is set incorrectly or the connected switch port is not configured accordingly, the interface can become unreachable.</p> <p>Default VLAN ID cannot be set if IPv6 is enabled.</p> <p>Default VLAN ID is supported in the clustered environment.</p> <p>In case of VMware ESXi based Virtual Appliance(VA), set the vSwitch configuration to port 4095 to allow ICS to tag the traffic.</p> <p>The set default VLAN ID should be added as a member in the physical port of switch and the same VLAN should be removed from native VLAN ID.</p>

Configuring the External Port

The external port connects to the Internet. You can use the System > Network pages to configure the external port.

To configure the external port:

1. Select **System > Network > External Port > Settings** to display the configuration page.

The following figure shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration as described in the following table.
3. Save your changes.

The following figure depicts the Ivanti Connect Secure External Port Configuration Page:

Network > External Port > Settings

Settings

Network Settings (for node node-1)

External Port - Settings

Settings for: node-1 (this node) Update

Overview Internal Port **External Port** Management Port VLANs Routes Hosts VPN Tunneling Proxy Server

Settings Virtual Ports ARP Cache ND Cache

Enter the network settings and click the Save Changes button at the bottom of the page.

▼ Use Port

Enabled Disabled

▼ IPv4 Settings

Enable IPv4 Disable IPv4

Note: Changing above setting might restart some services. This restart might drop all the connections to the Pulse Connect Secure.

*IP Address:

*Netmask:

*Default Gateway:

Note: If you need to specify static routes, you can do so on the [Static Routes](#) page.

▼ IPv6 Settings

Enable IPv6 Disable IPv6

Note: Changing above setting might restart some services. This restart might drop all the connections to the Pulse Connect Secure.

Link Local Address:

*IPv6 Address:

*Prefix Length: (1 to 128)

*Default Gateway:

▼ Advanced Port

MAC Address:

Link Speed:

*ARP Ping Timeout: seconds (3 to 300 seconds)

*MTU: bytes (Maximum Transmission Unit. If IPv6 is enabled (1280 to 1500), else (576 to 1500).)

Default VLAN ID:

Default VLAN ID for the traffic for this port.
WARNING: Adding VLAN ID may break connectivity unless corresponding switch port is configured to handle tagged traffic.

Save Changes

The following table lists the External Port Configuration Guidelines:

Settings	Guidelines
Use Port?	
Use Port?	Select Enabled to use the port; otherwise, select Disabled.
IPv4 Settings	
IP Address	Specify an IP address. An IP address is an identifier for a computer or device on a TCP/IP network. Networks using the TCP/IP protocol route messages based on the IP address of the destination. The format of an IPv4 address is a 32-bit numeric address written as four numbers separated by periods. Each number can be 0 to 255.
Netmask	Specify a netmask. A netmask indicates which part of an IP address indicates network identification and which part indicates the host identification. For example, the IP address and netmask 10.20.30.1 255.255.255.0 (or 10.20.30.1/24) refer to all the hosts in the 10.20.30.0 subnet. The IP address and netmask 10.20.30.1 255.255.255.255 (or 10.20.30.1/32) refer to a single host.
Default Gateway	Specify the IPv4 address for the default gateway for the routing domain to which the device belongs. A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.
IPv6 Settings	
Enable IPv6 / Disable IPv6	Disabled by default. Enable to support access from IPv6 endpoints. When you enable IPv6, the system acquires a link local address. If you switch from enabled to disabled, the system clears the link local address.
Link Local Address	Display the autoconfigured link local address (after you have enabled and saved the IPv6 configuration).
IPv6 Address	Specify a routable IPv6 address, such as a global unicast address that your network plan has provisioned for this host and interface. Automatic configuration methods are not supported. You must specify the appropriate address manually.
Prefix Length	Specify how many of the higher order contiguous bits of the IPv6 address comprise the prefix (the network portion of the IPv6 address). The default is 64.

Settings	Guidelines
Gateway	<p>Specify the IPv6 address for the default gateway for the routing domain to which the device belongs.</p> <p>A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.</p>
Advanced Settings	
MAC Ad-dress	Display the MAC address for the interface.
Link Speed	<p>Specify the speed and duplex combination for the interface.</p> <p>If you run SNMP_GET and then change the Link Speed value, you must wait at least 5 minutes after submitting the change before running SNMP_GET again.</p>
ARP Ping Timeout	<p>(IPv4 only.) Specify how long the system should wait for responses to Address Resolution Protocol (ARP) requests before timing out. Cluster nodes send ARP requests to the gateways of other nodes to determine if they are properly communicating with one another.</p> <p>If you have not deployed a cluster, the system does not use this setting. If the node belongs to a cluster, the timeout interval that you specify is synchronized across the cluster. In multisite clusters, you can override this setting for the individual nodes in the cluster using options in the System > Clustering page. Use caution when changing this setting in active/passive clusters, however, because the system also uses the ARP Ping Timeout setting on the Internal tab as a fail-over timer for the VIP.</p>
MTU	<p>Specify the maximum transmission unit.</p> <p>If IPv6 is enabled, the valid range is 1280 to 1500. If IPv6 is not enabled, the valid range is 576 to 1500.</p> <p>We recommend you retain the default MTU setting (1500) unless you must change the setting for troubleshooting purposes.</p> <p>If the administrator sets ignore-tcp-mss in Advanced Client Configuration, then the TCP MSS option is ignored during the virtual adapter MTU calculation on the client side. For details, see Using the Advanced Client Configuration Feature</p>

Settings	Guidelines
Default VLAN ID	<p>(Optional) Specify the default VLAN ID for the traffic of this port. When this parameter is set, all the traffic on this interface is subsequently tagged with the set VLAN ID and also accepts only incoming traffic with the same tag. Necessary changes are required on the connected switch port to handle bi-directional tagged traffic.</p> <p>If default VLAN ID is set incorrectly or the connected switch port is not configured accordingly, the interface can become unreachable.</p> <p>Default VLAN ID cannot be set if IPv6 is enabled.</p> <p>Default VLAN ID is not supported in a clustered environment.</p> <p>In case of VMware ESXi based Virtual Appliance(VA), set the vSwitch configuration to port 4095 to allow ICS to tag the traffic.</p> <p>The set default VLAN ID should be added as a member in the physical port of switch and the same VLAN should be removed from native VLAN ID.</p> <p>If default VLAN ID is set incorrectly or the connected switch port is not configured accordingly, the interface can become unreachable.</p>

Using the Internal and External Ports

The internal port, also known as the internal interface, handles all LAN requests to resources, listening for Web browsing, file browsing, authentication, and outbound mail requests. You configure the internal port by providing IP address, gateway, DNS server and domain, and MTU settings during the initial setup of Ivanti Connect Secure. You can also change them on the System > Network > Internal Port > Settings tab. Alternatively, you can deploy the appliance in dual-port mode to listen for incoming Web and mail proxy SSL connections on an external port.

The external port, also known as the external interface, handles all requests from users signed into Ivanti Connect Secure from outside the customer LAN, for example, from the Internet. Before sending a packet, Ivanti Connect Secure determines if the packet is associated with a TCP connection that was initiated by a user through the external interface. If that is the case, Ivanti Connect Secure sends the packet to the external interface. All other packets go to the internal interface.

The routes that you specify for each interface apply after Ivanti Connect Secure has determined whether to use the internal or external interface. No requests are initiated by Ivanti Connect Secure from the external interface, and this interface does not accept any other connections (except ping and traceroute connections). All requests to any resource are issued from the internal interface.



If you enable the external port, then, by default, administrators may no longer sign in from an external location. You can open the external port for administrators on the Administrators > Admin Realms > RealmName > Authentication Policy > Source IP tab.

Using the Management Port

This topic describes how to configure the management port.

Management Port Overview

You connect the management port to an Ethernet switch or router that is part of your internal local area network (LAN) and that can connect to your network management infrastructure. When the management port is enabled, the following traffic is directed out the management port: archiving (FTP/SCP), NTP, push config, SNMP, syslog. When the management port is not enabled, that traffic uses the internal port.

On Policy Secure systems, you cannot configure the user realm configuration, the RA-DIUS client configuration, or the Infranet Enforcer configuration to use the management port.

Supported Platforms

The following hardware platforms are equipped with a management port:


- PSA300, PSA3000, PSA5000
- PSA7000c and PSA7000f

Configuring the Management Port

To configure the management port:

1. Select **System > Network > Management Port > Settings** to display the configuration page. The following figure shows the configuration page for Ivanti Connect Secure.
2. Complete the configuration as described in the following table.
3. Save your changes.

The following figure depicts the Ivanti Connect Secure Management Port Configuration Page:



The following table lists the Management Port Configuration Guidelines

Settings	Guidelines
Use Port?	
Use Port?	Select Enabled to use the port; otherwise, select Disabled.
IPv4 Settings	
IP Address	Specify an IP address. An IP address is an identifier for a computer or device on a TCP/IP network. Networks using the TCP/IP protocol route messages based on the IP address of the destination. The format of an IPv4 address is a 32-bit numeric address written as four numbers separated by periods. Each number can be 0 to 255.

Settings	Guidelines
Netmask	A netmask indicates which part of an IP address indicates network identification and which part indicates the host identification. For example, the IP address and netmask 10.20.30.1 255.255.255.0 (or 10.20.30.1/24) refer to all the hosts in the 10.20.30.0 subnet. The IP address and netmask 10.20.30.1 255.255.255.255 (or 10.20.30.1/32) refer to a single host.
Default Gateway	Specify the IPv4 address for the default gateway for the routing domain to which the device belongs. A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.
IPv6 Settings	
Enable IPv6 / Disable IPv6	Disabled by default. Enable to support network management traffic over IPv6 networks. When you enable IPv6, the system acquires a link local address. If you switch from enabled to disabled, the system clears the link local address.
Link Local Address	Display the autoconfigured link local address (after you have enabled and saved the IPv6 configuration).
IPv6 Ad-dress	Specify a routable IPv6 address, such as a global unicast address that your network plan has provisioned for this host and interface. Automatic configuration methods are not supported. You must specify the appropriate address manually.
Prefix Length	Specify how many of the higher-order contiguous bits of the IPv6 address com-prise the prefix (the network portion of the IPv6 address). The default is 64.
Gateway	Specify the IPv6 address for the default gateway for the routing domain to which the device belongs. A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.
Advanced Settings	
MAC Ad-dress	Display the MAC address for the interface.

Settings	Guidelines
Link Speed	Specify the speed and duplex combination for the interface. If you run SNMP_GET and then change the Link Speed value, you must wait at least 5 minutes after submitting the change before running SNMP_GET again.
ARP Ping Timeout	(IPv4 only.) Specify how long the system should wait for responses to Address Resolution Protocol (ARP) requests before timing out. Cluster nodes send ARP requests to the gateways of other nodes to determine if they are properly communicating with one another. If you have not deployed a cluster, the system does not use this setting. If the node belongs to a cluster, the timeout interval that you specify is synchronized across the cluster. In multisite clusters, you can override this setting for the individual nodes in the cluster using options in the System > Clustering page. Use caution when changing this setting in active/passive clusters, however, because the system also uses the ARP Ping Timeout setting on the Internal tab as a fail-over timer for the VIP.
MTU	Specify the maximum transmission unit. If IPv6 is enabled, the valid range is 1280 to 1500. If IPv6 is not enabled, the valid range is 576 to 1500. We recommend you retain the default MTU setting (1500) unless you must change the setting for troubleshooting purposes.
Default VLAN ID	(Optional) Specify the default VLAN ID for the traffic of this port. When this parameter is set, all the traffic on this interface is subsequently tagged with the set VLAN ID and also accepts only incoming traffic with the same tag. Necessary changes are required on the connected switch port to handle bi-directional tagged traffic. If default VLAN ID is set incorrectly or the connected switch port is not configured accordingly, the interface can become unreachable. Default VLAN ID cannot be set if IPv6 is enabled. Default VLAN ID is not supported in a clustered environment. In case of VMware ESXi based Virtual Appliance (VA), set the vSwitch configuration to port 4095 to allow ICS to tag the traffic. The set default VLAN ID should be added as a member in the physical port of switch and the same VLAN should be removed from native VLAN ID.

Using the Serial Console to Configure the Management Port

To configure management port network settings from the serial console:

1. Start a serial console session.
2. Select item **1, System Settings and Tools**.
3. Select item **10, Configure Management port**. The text indicates if the option is enabled or disabled.
4. Enter the network settings for the Management Port, as prompted.



If you enable the Management Port but neglect to configure the IP address and net-mask, the port reverts to a disabled state. Also, you cannot clear Management Port settings from the serial console when the port is disabled, though you can clear them from within the admin console.

5. When prompted to accept the changes, if they are correct, enter y. Otherwise, repeat the process to correct the settings.
6. Close the serial console.

Configuring Administrator Access

You can configure the Administrators > Admin Realm > Authentication Policy > Source IP restrictions configuration to enable administrator sign-in through the management port.

You can use Administrator realms to control administrator access to system ports, including the management port.

To control administrator access to the management port:

1. Enable the management port.
2. Perform one of the following steps:
 - Select **Administrators > Admin Realms > Admin Users** to modify the default admin users realm.
 - Select **Administrators > Admin Realms**, then click **New**, to create a new administrator realm.
3. Select the **Authentication Policy > Source IP**.
4. Select one of the following options:

- Allow users to sign in from any IP address-Allows users to sign in from any IP address to satisfy the access management requirement.
- Allow or deny users from the following IP addresses-Specifies whether to allow or deny users access from all of the listed IP addresses, based on their settings.

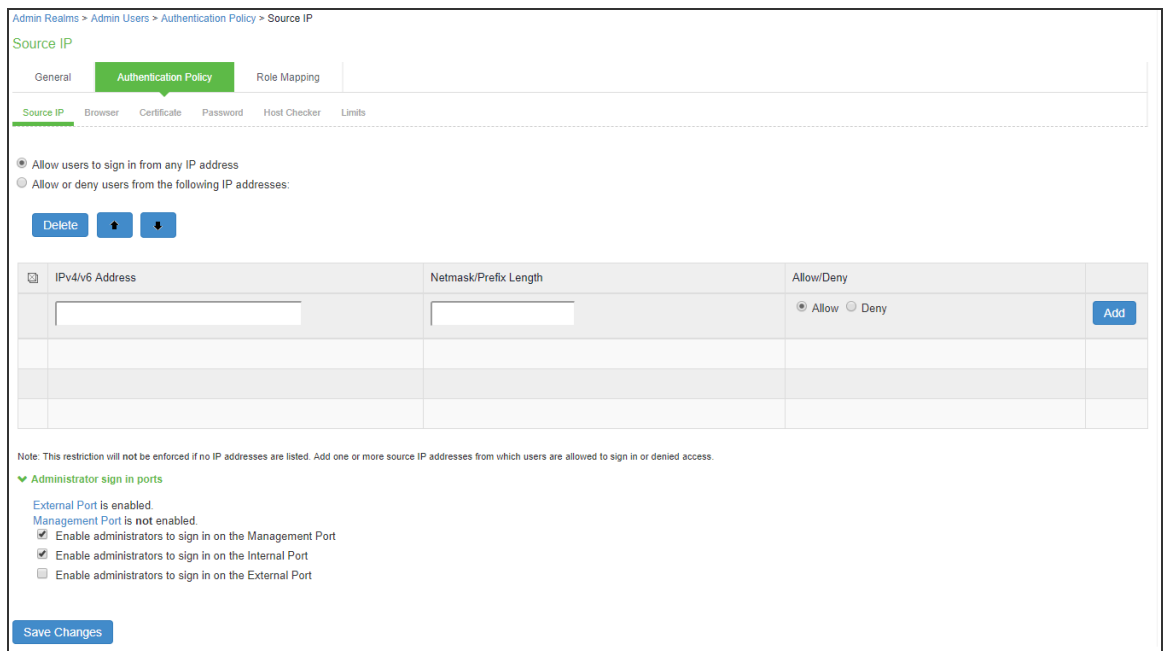
To specify access from an IP address:

- Enter the IP address and netmask.
 - Select either Allow to allow users to sign in from the specified IP address, or Deny to prevent users from signing in from the specified IP address.
5. Select the available options to allow administrators to sign in to all available ports, to the management port or the internal port only, or to restrict them from signing in to any of the ports. In some cases, you may inadvertently limit administrative access completely. If this occurs, you can reconfigure the ports by way of the serial console.

Select from the following available options:

- Enable administrators to sign in on the management port.
- Enable administrators to sign in on the internal port.
- Enable administrators to sign in on the external port.

The following figure shows the configuration page for administrator access.



6. Click **Save Changes**.

Configuring VLAN Ports

Your network design might include VLANs to provide network segmentation. When connected to a trunk port on a VLAN-enabled switch, the system encounters traffic from all VLANs. This is useful for network designs with separate VLANs for separate classes of users or endpoints, and for making the system accessible from all VLANs. You can use RADIUS attributes to place different users in different network segments.

The system supports IEEE 802.1Q VLAN tagging. You must define a VLAN port for each VLAN. The internal port must be assigned to the root system and must be marked as the default VLAN. Routes to servers reachable from the VLAN interfaces must have the next-hop gateway set to the configured gateway for the VLAN interface, and must have the output port defined as the VLAN port.

When you save the configuration for a new VLAN port, the system creates two static routes by de-fault:

- The default route for the VLAN pointing to the default gateway.
- The interface route to the directly connected network.


To configure an internal VLAN port:

1. Select **System > Network > VLANs > Internal Port > New VLAN Port -Settings**.

The following figure shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration as described in the following table.
3. Save your changes.

The following figure depicts the Ivanti Connect Secure VLAN Port Configuration Page:



The following table lists the VLAN Port Configuration Guidelines

Settings	Guidelines
Use Port?	
Use Port?	Select Enabled to use the port; otherwise, select Disabled.
VLAN Settings	
Port Name	Specify a name that is unique across all VLAN ports that you define on the system or cluster. Only alphanumeric characters, "-", or "_" are allowed.
VLAN ID	Specify a number between 1 and 4094. The VLAN ID assignment must be unique on the system.
IPv4 Settings	

Settings	Guidelines
IP Address	<p>Specify an IP address and netmask combination that is from the same network as the VLAN. VLAN IP addresses must be unique. You cannot configure a VLAN to have the same network as the internal port. For example, if the internal port is 10.64.4.30/16 and you configure a VLAN as 10.64.3.30/16, you might get unpredictable results and errors.</p> <p>The format of an IPv4 address is a 32-bit numeric address written as four numbers separated by periods. Each number can be 0 to 255.</p>
Netmask	<p>Specify a netmask. A netmask indicates which part of an IP address indicates network identification and which part indicates the host identification. For example, the IP address and netmask 10.20.30.1 255.255.255.0 (or 10.20.30.1/24) refer to all the hosts in the 10.20.30.0 subnet. The IP address and netmask 10.20.30.1 255.255.255.255 (or 10.20.30.1/32) refer to a single host.</p>
Default Gateway	<p>Specify the IPv4 address for the default gateway for the routing domain to which the device belongs.</p> <p>A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.</p>
IPv6 Settings	
IPv6 Settings	Select Enabled to use the port; otherwise, select Disabled .
IPv6 Address	Specify a routable IPv6 address, such as a global unicast address that your network plan has provisioned for this host and interface. Automatic configuration methods are not supported. You must specify the appropriate address manually.
Prefix Length	Specify how many of the higher order contiguous bits of the IPv6 address comprise the prefix (the network portion of the IPv6 address). The default is 64.
Default Gateway	<p>Specify the IPv6 address for the default gateway for the routing domain to which the device belongs.</p> <p>A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.</p>

-
- Link speed, ARP ping timeout, and MTU settings are inherited from the internal port configuration.



- To configure an external VLAN port, Select **System > Network > VLANs > External Port > New VLAN Port -Settings**.
 - To configure a Management port, Select **System > Network > VLANs > Management Port > New VLAN Port -Settings**.
-

Then, complete the configuration as described in the VLAN Port Configuration Guidelines table in the next section.

Using Virtual Ports

This topic describes virtual ports.

Configuring Virtual Ports

You can use virtual ports to provide different groups of users access to the same system using different IP aliases and domains.

Virtual ports are associated with the physical internal port and physical external port. The virtual port shares all of the network settings with the associated physical port, except for the IP address.

When you configure virtual ports, you in essence are creating name-IP address pairs. The names and IP addresses must be unique in your network. An alias can include IPv4 addresses, IPv6 addresses, or both. However, the corresponding IP protocol must be enabled on the physical port for the addresses to take effect.

To configure a virtual port:

1. Select **System > Network > PortName > Virtual Ports**. *PortName* is Internal Port or External Port.
2. Click **New Port** to display the configuration page.

The following figure shows the configuration page for Ivanti Connect Secure.

3. Complete the configuration as described in the following table.
4. Save your changes.

The following figure depicts the Ivanti Connect Secure Virtual Port Configuration Page:

The following table lists the Virtual Port Configuration Guidelines:

Settings	Guidelines
Name	Specify a name for the virtual port. The names and IP addresses in the virtual port configuration must be unique in your network.
Physical Port	Display the name of the physical port associated with the virtual port. The virtual port inherits link speed, ARP ping timeout, and MTU settings from the physical port configuration.
IPv4 Ad-dress	Specify an IPv4 address. An alias can include IPv4 addresses, IPv6 addresses, or both. However, the corresponding IP protocol must be enabled on the physical port for the addresses to take effect.
IPv6 Ad-dress	Specify an IPv6 address. An alias can include IPv4 addresses, IPv6 addresses, or both. However, the corresponding IP protocol must be enabled on the physical port for the addresses to take effect.

Using Device Certificates with Virtual Ports

Virtual ports can be used to create multiple fully qualified domain names for user sign-in. When a user tries to sign in using the IP address defined in a virtual port, the system presents the certificate associated with the virtual port to initiate the SSL transaction.

You can approach the digital certificate security and virtual ports implementation in either of the following ways:

- Associate all hostnames with a single certificate-With this approach, you use a single wildcard certificate to validate the identity of all system hostnames, regardless of which hostname is used to sign in. A wildcard certificate includes a variable element in the domain name, making it possible for users who sign in from multiple hosts to map to the "same" domain. For example, if you create a wild-card certificate for *.yourcompany.com, the system uses the same certificate to validate its identity to users who sign in to employees.yourcompany.com as it does to users who sign into part-ners.yourcompany.com.
- Associate each hostname with its own certificate-With this approach, you associate different hostnames with different certificates. Create a virtual port for each hostname. A virtual port activates an IP alias on a physical port. For example, you can create two virtual ports on a single appliance, map-ping the first virtual port to the IP address 10.10.10.1 (sales.yourcompany.com) and the second virtual port to the IP address 10.10.10.2 (partners.yourcompany.com). Then you can associate each of these virtual ports with its own certificate, ensuring that users authenticate through different certificates.

To associate certificates with virtual ports:

1. Create virtual ports.
2. Import the device certificates.

3. Associate the device certificates with the virtual ports:

1. Select **System > Configuration > Certificates > Device Certificates**.
2. Click the link of the device certificate you want to configure to display the configuration page.

The following figure shows the configuration page for Ivanti Connect Secure.

3. Use the controls in the "Present certificate on these ports" section to associate ports with the certifi-cate.

The following figure depicts the Ivanti Connect Secure Certificate Details Page:

Certificate Details

▼ Certificate

Issued To: sdklcsdn.psecure.net
Issued By: ??
Valid: Dec 21 10:17:35 2015 GMT to Jun 12 10:17:35 2021 GMT
Details: ▼ Other Certificate Details

Version: 1
Serial: 9f:da:b4:82:0a:2f:da:d1
Signature Algorithm: sha1WithRSAEncryption
Public Key Algorithm: rsaEncryption
Public Key Type: RSA
Public Key Bits: 2048
Public Key:

```
Public-Key: (2048 bit)
Modulus:
 00:cc:f7:7a:a7:7f:3b:2f:46:0b:61:e7:45:76:b9:
 3c:db:46:0a:72:4e:a6:e7:0f:28:4d:4d:19:aa:5f:
 11:e0:a0:97:08:b4:42:9d:42:dd:c5:a3:0f:bb:06:
 68:81:5e:da:52:46:e8:09:e7:ba:23:46:89:15:dd:
 79:bd:d5:7d:ed:83:2e:4f:1a:b6:30:c0:d8:32:5b:
 2d:52:09:a5:cd:19:2f:7d:91:b8:2e:bb:33:b9:fd:
 38:a8:f2:0a:60:c6:e1:eb:70:f3:88:e4:85:cc:88:
 ce:d0:c5:b4:60:82:48:1e:28:e2:d8:d3:b4:3c:83:
 87:37:fc:36:9e:da:24:c7:5f:b9:cb:80:7a:a8:6c:
 5b:53:a4:a4:46:dd:03:3e:7e:e7:8c:ba:4e:0d:cd:
 cd:46:aa:df:b0:df:a5:e8:b0:64:76:02:ee:d2:6f:
 1d:91:3e:79:80:42:76:30:5a:b1:54:38:07:dd:9a:
 d0:06:10:60:b9:14:55:f5:c6:39:19:fc:31:d9:9e:
 e7:63:ad:7b:61:57:2f:24:1d:ab:ce:f1:4a:55:18:
 f5:1f:fe:93:c3:86:02:9b:35:06:fd:a8:4c:12:75:
 d0:dd:06:8b:03:41:09:05:2c:b3:1b:1b:8d:7c:28:
 55:ef:3c:f8:5c:57:af:e3:d2:c0:79:1c:21:87:1c:
 42:69
Exponent: 65537 (0x10001)
```

Thumbprint Algorithm: SHA1
Thumbprint: A8:BA:55:09:6A:3D:E2:FD:37:8C:AF:44:AB:00:0D:66:EB:84:6B:7E

Download

▼ Present certificate on these ports

Select the internal and external virtual ports that will present this certificate:

Internal Virtual Ports: **Selected Virtual Ports:**

<Internal Port>

Add ->

Remove

<Internal Port>

External Virtual Ports: **Selected Virtual Ports:**

<External Port>

Add ->

Remove

<External Port>

Vlan Ports: **Selected Vlan Ports:**

<Vlan Port>

Add ->

Remove

<Vlan Port>

Configuring the System Date and Time

You can use the admin console to set the system date and time manually or by configuring a network time protocol (NTP) server. The system supports NTPv4, which is backwards compatible with NTPv3 and NTPv2.



BEST PRACTICE: We recommend you use NTP to synchronize the date and time clocks on all network systems. Using NTP obviates issues that might occur with cluster synchronization, network communication that uses time-sensitive protocols, such as SAML, and implementation of time-based policies, such as local authentication server account expiration. In addition, using NTP as a standard in your network rationalizes timestamps in logs, which facilitates reporting and troubleshooting.

On a VMware virtual appliance, the cockpit data may be erased each hour if the same NTP server is not defined here, on the Ivanti Connect Secure license server, and on the ESXi server.

To set the system date and time:

1. Select **System > Status > Overview** to display the System Status dashboard.
2. Click the **System Date and Time Edit** link to display the configuration page.

The following figure depicts the NTP:

Status > Overview > Date and Time

Date and Time

System Date: 7/18/2020
System Time: 10:09:54 AM

Time Zone: (GMT+05:30) Kolkata, Chennai, Mumbai, New Delhi

▼ Time Source

Use Pool of NTP Servers

Configure pool of NTP servers (IP Address/Hostname)
Please make sure NTP server is reachable via port configured at [Advanced Networking](#) page.
For troubleshooting use ntpq command under [Troubleshooting](#) page

* NTP Server 1:	<input type="text" value="time.pulsesecure.net"/>	Key 1:	<input type="text" value="•••••"/>	(optional)
NTP Server 2:	<input type="text" value="m10.96.195.10"/>	Key 2:	<input type="text"/>	(optional)
NTP Server 3:	<input type="text" value="0.pool.ntp.org"/>	Key 3:	<input type="text" value="•••••"/>	(optional)
NTP Server 4:	<input type="text" value="1.pool.ntp.org"/>	Key 4:	<input type="text"/>	(optional)


Set Time Manually

Date: (mm/dd/yyyy)

Time: AM (hh:mm:ss)


For troubleshooting, navigate to **Maintenance > Troubleshooting > Tools > Commands** and then use **ntpq** command.

The following figure depicts the ntpq Command:



3. Complete the configuration as described in the following table.
4. Save the configuration.

The following table lists the Date and Time Configuration Guidelines:

Settings	Guidelines
Time Zone	Select your time zone. Selecting the appropriate time zone enables the system to automatically adjust the time for Daylight Saving Time changes.
Time Source	
Use Pool of NTP Servers	<p>Select this option to configure pool of NTP servers. Configuring one NTP server is mandatory and keys are optional.</p> <hr/> <p> ICS VMs deployed on VMWare ESX server will synchronize time with ESXi host. To use NTP/local time, turn off VMWare Tools Time Synchronization completely.</p> <hr/> <p>BEST PRACTICE: It is not recommended to use only two NTP servers. If more than one NTP server is required, four NTP servers is recommended minimum. Four servers protects against one incorrect timesource.</p>
NTP Server(s)	Specify the fully qualified domain name or IP address for the NTP server.

Settings	Guidelines
Key(s)	<p>If you are using NTPv4, specify the symmetric key. The key must be pre-synchronized with the NTP server. For example, if you want to configure NIST's clock as the NTP server, you must request a key beforehand and have NIST send that key to you.</p> <p>The key for MD5 is in the following format: KeyNumber M KeyValue The key for SHA1 is in the following format: KeyNumber SHA1KeyValue</p>
Set Time Manually	
Date	Specify the date. You can click Get from Browser to automatically populate the Date and Time fields.
Time	Specify the time and select AM or PM.

Configuring Network Services

You configure DNS and WINS services when you initially configure the system with the serial console. If necessary, you can use the System > Network > Overview page to modify the configuration. You can also use this page to configure a hostname.

The network services overview page also displays the node name (if the node belongs to a cluster), and the status and interface statistics for the internal port, external port, and management port.

To configure network services:

1. Select **System > Network > Overview** to display the configuration page.

The following figure shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration as described in the following table.

3. Save your changes.

The following figure depicts the Ivanti Connect Secure Network Services Configuration Page:

Network > Overview

Overview

Network Settings
Overview

Overview Internal Port External Port Management Port VLANs Routes Hosts VPN Tunneling Proxy Server

Enter the network settings and click the Save Changes button at the bottom of the page.

Status

Internal Port: Connected, Speed: 10000Mb/s, Duplex: full
 RxPacket: 732462 RxError: 0 RxDrop: 0 RxMulticast: 0
 TxPacket: 725518 TxError: 0 TxDrop: 0 TxMulticast: 0

External Port: Connected, Speed: 10000Mb/s, Duplex: full
 RxPacket: 2647415 RxError: 0 RxDrop: 0 RxMulticast: 206459
 TxPacket: 353214 TxError: 0 TxDrop: 0 TxMulticast: 0

Management Port: Connected, Speed: 10000Mb/s, Duplex: full
 RxPacket: 31290320 RxError: 0 RxDrop: 0 RxMulticast: 0
 TxPacket: 290012 TxError: 0 TxDrop: 0 TxMulticast: 0

Network Identity

Hostname: Fully-qualified hostname (example: device.domain.com)

DNS name resolution

Primary DNS: IPv4/IPv6 address

Secondary DNS: IPv4/IPv6 address

VPN tunnels maximum bandwidth: Mbps Maximum bandwidth for VPN tunnel traffic.

Note: These fields are required to use the bandwidth management feature. The Total Maximum Bandwidth must be greater than the VPN Tunnels Maximum Bandwidth. Set to 0 to disable this feature.

IPv6 Settings

Disable ICMPv6 echo response for multicast echo requests

Disable ICMPv6 destination unreachable response

DSCP value: 0 - 63 (Applied to the dscp field of all the IPv6 packets originated or forwarded from the device.)

Tunnel Settings

TCP MSS Value: Bytes TCP MSS value for tunnels between 0 and 1360. Set to 0 to disable this feature

The following table lists the Network Services Configuration Guidelines:

Settings	Guidelines
Status	
Status	Display node name, interface statistics for the internal port, external port, and management port.
Network Identity	
Hostname	Specify a fully qualified hostname. For example, domain.company.com. The hostname cannot exceed 30 characters
DNS Name Resolution	
Primary DNS	Specify the IP address for the primary DNS server.
Secondary DNS	Specify the IP address for the secondary DNS server.

Settings	Guidelines
DNS Domain(s)	Specify a comma-separated list of default domains. The system searches the domains in the order they are listed.
Preferred DNS Response	This field determines what DNS requests and responses ICS will prefer to the configured DNS server. Select 'V4' if ICS is interested only in IPv4 hostname resolution requests and responses to/from the backend DNS server. Select 'Both' if ICS needs to send and receive both IPv4 and IPv6 host-name resolution requests and responses.
Port for DNS Traffic	Prior to 9.1R1 release, DNS traffic was sent over the Internal interface. Starting with 9.1R1 release, an administrator can modify the DNS setting to any physical interface namely Internal Port, External Port or Management Port. In case of a fresh installation or an upgrade, DNS port will be set to Internal port. In case of a cluster, the setting can be made node-specific as well as cluster-wide.
Windows Networking	
WINS	Specify the hostname or IP address of a local or remote Windows Internet Naming Service (WINS) server that you use to associate workstation names and locations with IP addresses.
Bandwidth Management	This feature is available only on Ivanti Connect Secure.
Total Maximum Bandwidth	Specify the maximum bandwidth for all traffic.
VPN Tunnels Maximum Bandwidth	Specify the maximum bandwidth for VPN tunnel traffic. The value of total maximum bandwidth must be greater than the value of VPN tunnels maximum bandwidth
IPv6 Settings	
Disable ICMPv6 echo response for multicast echo requests	Allows enabling or disabling of the ability to send an Echo Reply message in response to an Echo Request message sent to an IPv6 multicast or anycast address.

Settings	Guidelines
Disable ICMPv6 destination unreachable response	Allows enabling or disabling the Destination Unreachable message in response to a packet that cannot be delivered to its destination for reasons other than congestion.
DSCP Value	Specify the value for verifying by packet capture at client side.
Tunnel Settings	
TCP MSS Value	Set the value of the MSS which can be ≤ 1460

Configuring NTP and Other Services Traffic Over Any Physical Interface

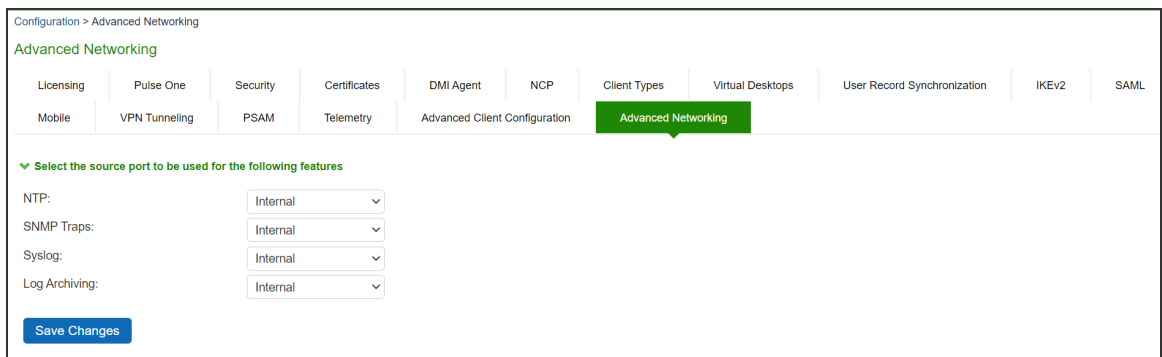
The NTP, SNMP, Syslog, and Log archiving services are set to send the traffic through Management port by default. In case the Management port is not available, the traffic is routed through Internal port. Now, an administrator can modify the settings of NTP and other services to any physical inter-face.

The following procedure describes the steps to configure the ports for the services. Before you proceed, ensure the External and Management ports are enabled for use in the network settings.

To configure Service Traffic Port Options:

1. Select **System > Configuration > Advanced Networking**.
2. For the individual service, select the required port from the drop-down list.

The following figure depicts the Source Port Selection:



In a cluster environment, when a node joins the cluster, configuration of the node is replaced with the configuration of other nodes in the cluster.

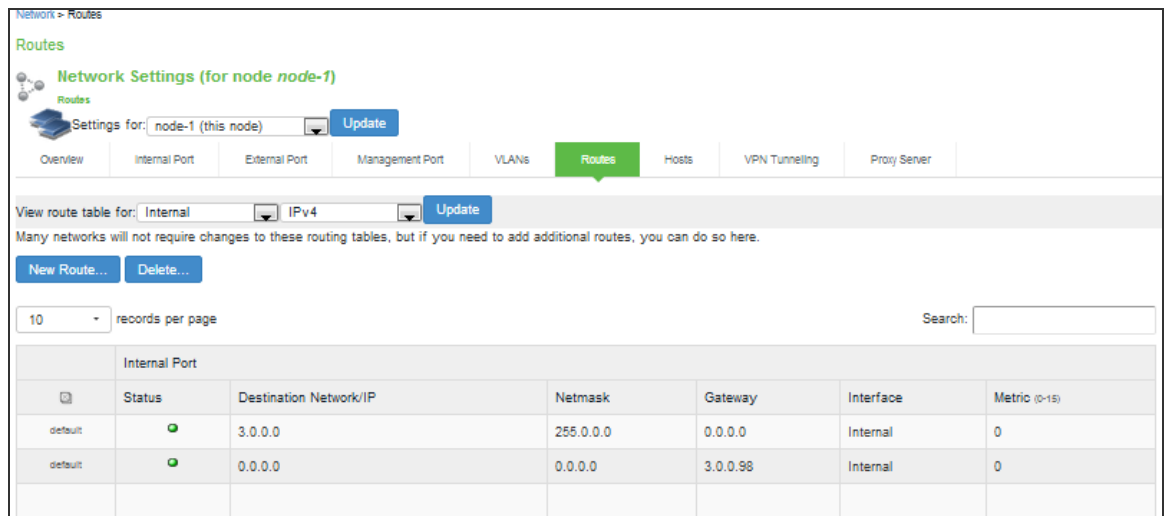
Managing the Routes Table

The system populates the routes table with dynamic, auto-discovered routes. Many networks will not require changes to this routing table. If necessary, you can delete routes or add static routes.

To manage the routes table:

1. Select **System > Network > Routes** to display the routes table.
shows the routes for Ivanti Connect Secure.
2. Use the controls described in the following table to manage the routes table.

The following figure depicts the Ivanti Connect Secure Routes Table:



The following table lists the Routes Table Controls

Controls	Description
View route table for	Use the controls to change the display to show the route table for internal, external, or management interfaces; and for IPv4 or IPv6 routes.
Delete	Select a row in the table and click Delete to delete a route.
New Route	Click New Route and complete the configuration to add a route to the table. You must specify a valid IP address, gateway, DNS address, and metric. The metric is a way of comparing multiple routes to establish precedence. Generally, the lower the number (from 0 to 15), the higher the precedence. Thus, a route with a metric of 2 is chosen over a route with a metric of 14.

Managing the Hosts Table

In general, the system uses the configured DNS servers to resolve hostnames, but it also maintains a local hosts table that can be used for name resolution. The system populates some entries from host-IP address pair settings in your configuration. You can add host-IP address mappings for other hosts that might not be known to the DNS servers used by the system, or in cases where DNS is not reachable.


To manage the hosts table:

Select **System > Network > Hosts** to display the hosts table.

The following figure shows the hosts table for Ivanti Connect Secure.

Use the controls described in the following table to manage the hosts table.

The following figure depicts the Ivanti Connect Secure Hosts Table:



The following table lists Hosts Table Controls:

Con-trols	Description
Add	Specify an IP address, hostname, and comment (a description for the benefit of sys-tem administrators) and click Add .
Delete	Click the delete icon in the last column to delete the row from the table.

Proxy Server Configuration

This feature provides communication between PSA-Vs with Pulse Cloud Licensing Server (PCLS) and Pulse One through a configured proxy server. A new tab called Proxy Server has been added in the Network Settings to configure the same.

To configure the proxy server settings:

1. Go to **System->Network->Proxy Server**.
2. Select the **Use Proxy Server for communicating with Pulse Cloud Licensing Service (PCLS)** check box.
3. Once enabled, the proxy server settings which include Host Name and Port must be set by the admin.
4. (Optional) If your proxy server requires further authentication, enter a username and password to log in to the proxy server.
5. Click on **Save**.

The following figure depicts the Proxy Server:

Network > Proxy Server

Proxy Server

Network Settings

Proxy Server

Overview Internal Port External Port Management Port VLANs Routes Hosts VPN Tunneling **Proxy Server**

Use this Proxy Server configuration during communication with following servers:

- Pulse Cloud Licensing Service (PCLS)
- Pulse One

Proxy Server

Preferred network interface:

Note: Please ensure that Preferred Network has IPv4 settings configured.

* Host Address: * Port:

Host can be a host name or a fully qualified domain name (e.g. "proxy.example.com") or an IPv4 address.

Username:

Password:

* indicates required field

Save



- If the global proxy server is configured and enabled for Pulse One, the local proxy settings configured in Pulse One is disabled. Similarly, if the global proxy server is configured and enabled for PCLS, the preferred network setting is disabled in the Download Licenses page.
- The **Proxy Server** tab is a cluster-wide setting for both active/active and active/passive clusters. Node-specific setting is disabled.

Managing the ARP Table

ARP stands for Address Resolution Protocol. In IPv4 networking, network nodes use ARP to maintain information about peer network nodes. ARP is used to associate the Layer 3 IP address with a Layer 2 MAC address of neighboring peer nodes. The system maintains an ARP table with dynamic, cached entries, and you can add static entries if necessary. The system caches dynamic entries for up to 20 minutes. Dynamic entries are deleted during a reboot. Static entries are restored after a reboot.


To manage the ARP table:

1. Select **System > Network > Port > ARP Cache**. *Port* is the Internal Port, External Port, or Management Port tab.

The following figure shows the ARP table for the internal Port for Ivanti Connect Secure.

2. Use the controls described in the following table to manage the ARP table.

The following figure depicts the Ivanti Connect Secure ARP Cache Table:



The following table lists the ARP Table Controls:

Controls	Description
Delete	Select a row in the table and click Delete to delete the entry.
Delete Dynamic Entries	Delete all dynamically discovered entries.
Add	Specify an IP address, a MAC address, and click Add to add an entry. If you add an entry that has the same IP address as an existing entry, the system over-writes the existing entry with your new entry. Also note that the system does not verify the validity of MAC addresses.

Managing the Neighbor Discovery Table

In IPv6 networking, network nodes use the **Neighbor Discovery Protocol (NDP)** to determine the Layer 2 MAC addresses for neighboring hosts and routers. The system uses NDP to maintain a cache of neighboring routers that are reachable and can forward packets on its behalf.

i In the current release, you can view discovered neighbors or clear the entire cache, but you cannot add neighbors or delete individual entries.

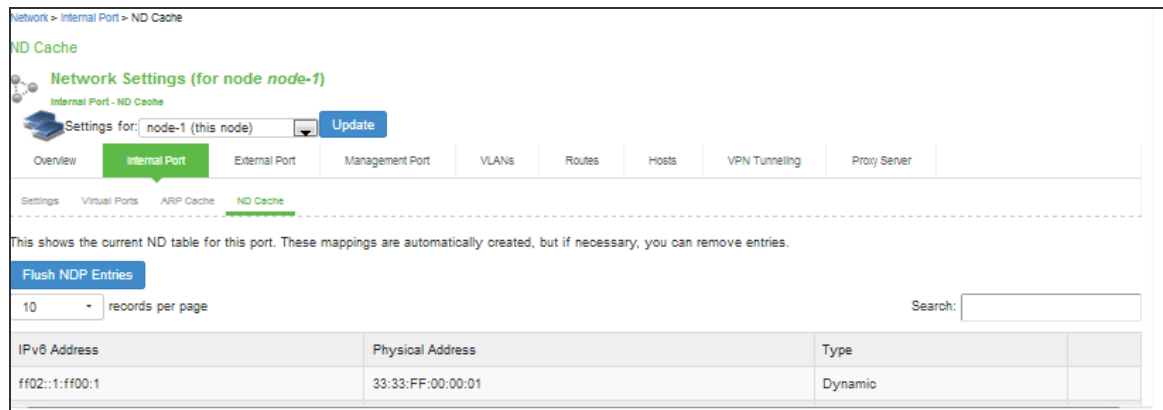
To manage the neighbor discovery table:

1. Select **System > Network > Port > ND Cache**. Port is the Internal Port, External Port, or Management Port tab.

The following figure shows the neighbor discovery table for the internal port for Ivanti Connect Secure.

2. Use the controls described in the following table to manage the neighbor discovery table.

The following figure depicts the Ivanti Connect Secure Neighbor Discovery Table:



The following table lists the Neighbor Discovery Table Controls

Controls	Description
Flush NDP Entries	Delete all dynamically discovered entries.

Using IPv6

This topic describes support for using IPv6.

Understanding IPv6

IP version 6 (IPv6) is an Internet Protocol designed to succeed IP version 4 (IPv4). This topic provides an overview of IPv6.

About IPv6

The ongoing expansive growth of the Internet and the need to provide IP addresses to accommodate it is escalating the emergent use of a new IP protocol. IPv6 was designed to satisfy the current and anticipated near future requirements.

IPv4 is widely used throughout the world today for the Internet, intranets, and private networks. IPv6 builds upon the functionality and structure of IPv4 in many aspects, including:

- Larger address space-IPv6 addresses are 128 bits long instead of 32 bits. This expands the address space from 4 billion addresses to over 300 trillion trillion addresses.
- New datagram format-The packet header is both simplified and enhanced to enable more secure and efficient routing.
- Improved fragmentation and reassembly-The maximum transmission unit (MTU) has been increased to 1280 bytes, for example.
- Transition mechanisms-Various network address translation (NAT) and tunneling mechanisms have been developed to support the transition to IPv6.

On February 3, 2011 Internet Assigned Numbers Authority (IANA) allotted the last block of IPv4 addresses to Regional Internet Registries (RIR), so the free pool of IPv4 addresses is now fully depleted. It is expected that in the near future Internet service providers (ISPs) will start issuing IPv6 addresses to their customers.

About IPv6 Address Types

[RFC 4291, IP Version 6 Addressing Architecture](#) describes the following types of IPv6 addresses:

- Unicast. An identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address.

- Anycast. An identifier for a set of interfaces. A packet sent to an anycast address is delivered to one of the interfaces identified by that address.
- Multicast. An identifier for a set of interfaces. A packet sent to a multicast address is delivered to all interfaces identified by that address.

The *link-local address* is a special IPv6 unicast address that is used in self-traffic and essential network communication, like Neighbor Discovery Protocol (NDP). When you enable IPv6 on a Ivanti Connect Secure interface, the system generates a link-local address that is derived from the interface MAC address.

When you configure IPv6 addresses for the system interfaces, you manually specify a routable address, such as global unicast address or an anycast address, depending on your routing implementation and your system deployment. A global unicast address must be globally unique so that it can be specified globally without need for modification. An anycast address represents a service rather than a specific device. An anycast address is not unique, but instead might be configured on each device in a cluster. You are not likely to use multicast addressing with Ivanti Connect Secure.

About IPv6 Address Text Representation

All IPv6 addresses are 128 bits long, written as 8 sections of 16 bits each. They are expressed in hexadecimal representation, so the sections range from 0 to FFFF. Sections are delimited by colons, and leading zeroes in each section may be omitted. If two or more consecutive sections have all zeroes, they can be collapsed to a double colon.

```
IPv6 addresses consist of 8 groups of 16-bit hexadecimal values separated by colons
(:). IPv6 addresses have the following format:
aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa
Each aaaa is a 16-bit hexadecimal value, and each a is a 4-bit hexadecimal value. The
following is a sample IPv6 address:
2001:0DB8:0000:0000:0008:0800:200C:417A
You can omit the leading zeros of each 16-bit group, as follows:
2001:DB8:0:0:8:800:200C:417A
You can compress 16-bit groups of zeros to double colons (::) as shown in the
following example, but only once per address:
2001:DB8::8:800:200C:417A
```

About the IPv6 Unspecified Address

In the IPv6 address space, the special "unspecified address" is 0:0:0:0:0:0:0:0. The compressed representation of the unspecified address is the double-colon (::). The unspecified address must never be assigned to a physical or virtual interface.

About the IPv6 Loopback Address

The special loopback address is the unicast address 0:0:0:0:0:0:0:1. The compressed representation of the loopback address is ::1. The loopback address may be used by a node to send an IPv6 packet to itself. It must not be assigned to a physical or virtual interface.

About IPv6 Address Prefixes

An IPv6 address prefix is a combination of an IPv6 prefix address and a prefix length used to represent a block of address space (or a network), similar to the use of an IPv4 subnet address and netmask combination to specify a subnet. An IPv6 address prefix takes the form ipv6-prefix/prefix-length. The ipv6-prefix variable follows general IPv6 addressing rules. The /prefix-length variable is a decimal value that indicates the number of contiguous, higher-order bits of the address that make up the network portion of the address. For example, 2001:DB8::/32 is an IPv6 address prefix, indicating that the first 32 bits make up the network portion of the address.

System Normalization of IPv6 Addresses

The system validates and normalizes IPv6 addresses entered by administrators. The normalized address is the address processed by the system, and it is the address that appears in logs.

The following table gives examples of how the system normalizes IPv6 address entries.

Example Entry	Normalized Address	Explanation
2001:DB8:1:1::3	2001:DB8:1:1::3	An address specified in compressed format is validated; the system uses the compressed form as the normalized form.
0:0:0::122	::122	Address is validated and normalized to compressed format.
FF01:0:0:0:0:0:101	FF01::101	Address is validated and normalized to compressed format.
2001:DB8::10.204.50.122	2001:DB8::ACC:327A	Address is validated and normalized to hexadecimal representation.
::FFFF:10.204.50.122	::FFFF:10.204.50.122	An address specified in compressed format is validated; the system uses the compressed form as the normalized form.

About Neighbor Discovery Protocol

[Neighbor discovery protocol \(NDP\)](#) allows different nodes on the same link to advertise their existence to their neighbors, and to learn about the existence of their neighbors.

Routers and hosts (nodes) use NDP messages to determine the link-layer addresses of neighbors that reside on attached links and to overwrite invalid cache entries. Hosts also use NDP to find neighboring routers that can forward packets on their behalf.

In addition, nodes use NDP to actively track the ability to reach neighbors. When a router (or the path to a router) fails, nodes actively search for alternatives to reach the destination.

IPv6 NDP corresponds to a number of the IPv4 protocols - ARP, ICMP Router Discovery, and ICMP Redirect. However, NDP provides many improvements over the IPv4 set of protocols. These improvements address the following:

- Router discovery-How a host locates routers residing on an attached link.
- Prefix discovery-How a host discovers address prefixes for destinations residing on an attached link. Nodes use prefixes to distinguish between destinations that reside on an attached link and those destinations that it can reach only through a router.
- Parameter discovery-How a node learns various parameters (link parameters or Internet parameters) that it places in outgoing packets.
- Address resolution-How a node uses only a destination IPv6 address to determine a link-layer address for destinations on an attached link.
- Next-hop determination-The algorithm that a node uses for mapping an IPv6 destination address into a neighbor IPv6 address (either the next router hop or the destination itself) to which it plans to send traffic for the destination.
- Neighbor unreachability detection-How a node determines that it can no longer reach a neighbor.
- Duplicate address detection-How a node determines whether an address is already in use by another node.

A router periodically multicasts a router advertisement from each of its multicast interfaces, announcing its availability. Hosts listen for these advertisements for address autoconfiguration and discovery of link-local addresses of the neighboring routers. When a host starts, it multicasts a router solicitation to ask for immediate advertisements.

The router discovery messages do not constitute a routing protocol. They enable hosts to discover the existence of neighboring routers, but they are not used to determine which router is best to reach a particular destination.

NDP uses the following Internet Control Message Protocol version 6 (ICMPv6) messages: router solicitation, router advertisement, neighbor solicitation, neighbor advertisement, and redirect.

NDP for IPv6 replaces the following IPv4 protocols: Router Discovery (RDISC), Address Resolution Protocol (ARP), and ICMPv4 redirect.

IPv6 Support Overview

This topic describes support for IP Version 6 (IPv6) networks.

Defining ESP Tunnel for Mixed Mode Traffic

To enable mixed mode traffic via ESP tunnel:

1. In the admin console, choose **System > Configuration > VPN Tunneling**.
2. In the IPv6 ESP Settings section, select the **Use ESP tunnel for 6in4 and 4in6 traffic** check box.
3. Click **Save Changes**.

To view the users connected via ESP tunnel, navigate to **System > Status > Active Users**.

Client Access Summary

Ivanti Connect Secure supports use of VPN Tunneling Connection Profile features to enable dual-stack endpoints to connect the Ivanti Connect Secure device and access corporate network IPv4 and IPv6 resources. The following table summarizes supported access scenarios. This is applicable to both SSL and ESP modes

The following table lists the Ivanti Connect Secure Client Access Scenarios

Endpoint	Connect Secure Interface	Tunnel	Re-source	Description of the Connection
IPv4/IPv6	IPv4	IPv4-in-IPv4	IPv4	All resource access policies are supported for access to IPv4 resources.
IPv4/IPv6	IPv6	IPv6-in-IPv4	IPv6	You must configure IPv4 and IPv6 address pools in the VPN Tunneling connection profile configuration. Access to IPv6 resources using VPN Tunneling connection profiles only.
		IPv4-in-IPv6	IPv4	You must configure IPv4 and IPv6 address pools in the VPN Tunneling connection profile configuration. All resource access policies are supported for access to IPv4 resources using VPN Tunneling connection profiles only.
		IPv6-in-IPv6	IPv6	You must configure IPv4 and IPv6 address pools in the VPN Tunneling connection profile configuration. All resource access policies are supported for access to IPv6 resources using VPN Tunneling connection profiles only.

The following table provides a summary of Pulse client and system software requirements for IPv6 deployment types.

Connect Secure	Pulse Client	IPv4-in-IPv4	IPv4-in-IPv6	IPv6-in-IPv4	IPv6-in-IPv6
9.1Rx	9.1Rx	Yes	Yes	Yes	Yes
9.0Rx	9.0Rx	Yes	Yes	Yes	Yes

Network Topologies


Ivanti Connect Secure release 8.0 and later supports Pulse client access to the IPv6 corporate network using VPN Tunneling Connection Profile features.

The role-based VPN Tunneling Connection Profile determines the IP addresses assigned to the client Pulse client virtual adapter. In this configuration, you must configure an IPv4 address pool. You configure an IPv6 address pool to enable access to IPv6 resources. When a client connects and is mapped to a role that includes the VPN Tunneling Connection configuration, the Pulse client virtual adapter is assigned all address from each pool-both an IPv4 and IPv6 address-and a single SSL tunnel is set up. When a connection is made to the system IPv4 address, the IPv4 traffic is encapsulated in the IPv4 tunnel ("4 in 4" tunneling), and the IPv6 traffic is encapsulated in the IPv4 tunnel ("6 in 4"). When a connection is made to the system IPv6 address, the IPv4 traffic is encapsulated in the IPv6 tunnel ("4 in 6"), and the IPv6 traffic is encapsulated in the IPv6 tunnel ("6 in 6").

In this release, the DNS server used by the system must be reachable by IPv4 and must be able to resolve both A and AAAA DNS queries. Only the VPN Tunneling Connection Profile is supported for access to IPv6 resources. All other connection options and resource policies are not supported for access to IPv6 resources.


The following figure shows a deployment topology for dual-stack-enabled endpoints that access the system over an ISP IPv4 network.

The following figure depicts the Dual Stack Endpoint Access Over ISP IPv4 Network:



The following figure shows a deployment topology for dual-stack-enabled endpoints that access the system over an ISP IPv6 network.

Dual Stack Endpoint Access Over ISP IPv6 Network:



IPv6 Support and Limitations for Ivanti Connect Secure Features

The following table summarizes IPv6 support and limitations for Ivanti Connect Secure features for Release 8.0 and later.

The following table lists the Summary of IPv6 Support

Feature	Summary
Pulse client access	<p>Only the Pulse client supports IPv6. The following behavior is expected for this release:</p> <p>Endpoints must have dual-stack enabled in order to access IPv6 resources over IPv4 networks.</p> <p>VPN Tunneling Connection Profiles support IPv4 and IPv6 address pools.</p> <p>VPN Tunneling Connection Profiles do not support ESP mode for IPv6 resource access. If a connection is configured for ESP mode, it automatically falls back to use SSL mode.</p> <p>On dual-stack endpoints, VPN Tunneling split tunneling rules are supported for both IPv4 and IPv6 based routes. The IPv4/IPv6 traffic allowed by a split tunneling policy is forwarded to the system in an IPv4/IPv6 tunnel.</p> <p>Legacy JSAM does not support IPv6.</p> <p>Pulse clients on the following platforms support VPN Tunneling connections for IPv6 resource access:</p> <p>Windows 8 (32 and 64 bit), Windows 10 Redstone</p> <p>Mac OS/X Snow Leopard, Lion, Mountain Lion, High Sierra, Mojave, Catalina</p> <p>Host Checker supports IPv6. Third-party Host Checker functionality is supported to the extent that it is IPv6-capable. For example, the following third-party components might require endpoints to connect over IPv4:</p> <p>Downloading antivirus signature updates from third-party vendors.</p> <p>Downloading Windows Patches from Microsoft download servers.</p>
Authentication	<p>Active Directory (Standard Mode) - IPv4 and IPv6 based Backend servers are supported.</p> <p>Radius Auth Server - IPv4 and IPv6 based Backend servers are supported.</p>
DNS	Supports both IPv4 and IPv6 DNS servers.
Administrator and management access	<p>The internal interface and management interface can be configured with an IPv4 address or dual stack (IPv4 and IPv6). The internal interface and management interface cannot be configured with only an IPv6 address because the system uses IPv4 for the connections with network services, including AAA, DHCP, and DNS.</p> <p>Typically, administrators access the administrator GUI through the internal interface or management interface, but you may enable administrator access through the external interface on the Authentication > Admin Realms > Admin Users > Authentication Policy > Source IP page.</p>

Feature	Summary
Configuration through the serial console	You cannot view or configure IPv6 network settings with the serial console.
External interface configuration	IPv4, IPv6, or both is supported.
Internal interface configuration	IPv4 or both IPv4 and IPv6 is supported. In other words, the internal interface must be configured for IPv4 connections; in addition, it may be configured for IPv6 connections. It may not be configured for IPv6 only.
Management interface configuration	IPv4 or both IPv4 and IPv6 is supported. In other words, the management interface must be configured for IPv4 connections; in addition, it may be configured for IPv6 connections. It may not be configured for IPv6 only.
Virtual interface configuration	An interface alias may include IPv4 addresses, IPv6 addresses, or both. However, the corresponding IP protocol must be enabled on the physical interface for the addresses to take effect.
VLAN configuration	IPv4, IPv6 or both is supported.
Clustering	Supports IPv6 configuration for active/active and active/passive clusters. The existing intra-cluster communication mechanism is preserved. The intra-cluster communication occurs over the IPv4 corporate network through the internal interfaces.
License server	IPv4 must be enabled for the "preferred network" you select for licensing protocol communication.
Web server	The implementation for IPv6 does not require reconfiguration of the system after upgrade. The Web server can listen for and accept IPv4 or IPv6 clients, and it can differentiate between them for internal purposes and for logging purposes.
ActiveSync	The implementation for IPv6 does not require reconfiguration of the system after upgrade. ActiveSync functionality is available to users connecting from IPv4 or IPv6 endpoints to an IPv4 backend server. Connection to an IPv6 backend server is not supported.

Feature	Summary
Connection profiles	<p>After upgrading, you can update your VPN Tunneling Connection Profile configuration to enable IPv6 address assignments to Pulse clients. You must configure a static IPv6 address pool. DHCPv6 is not supported.</p> <p>Also note that the IP address server configuration on the System > Network > VPN Tunneling page does not support filtering for IPv6 address pools. In active/active clusters, separate connection profiles need to be created with different IPv6 address pools for each node.</p> <p>WINS is not used in IPv6 networks; therefore, WINS settings are not applicable for connection profiles used for IPv6 access.</p> <p>The server-side proxy feature does not support IPv6.</p>
Resource policies	<p>You can configure VPN Tunneling Connection Profiles to enable access to all IPv6 resources in your corporate network; however, you cannot configure VPN Tunneling Access Control Policies to allow or deny access to particular IPv6 resources. As a workaround, we recommend you deploy firewall security to restrict access to IPv6 resources.</p> <p>To enable access to IPv6 resources, the DNS server used by the system must be reachable by IPv4 and must be able to resolve AAAA DNS queries.</p>
Core Access - Rewriter	
	<p>The implementation for IPv6 does not require reconfiguration of the system after upgrade. After upgrade, IPv6 endpoints can access internal IPv4 resources through the system. This applies to all system content rewriters: HTML, Java Script, Applets, VB Script, Flash, CSS, XML, PDF.</p> <p>You cannot configure Web Rewriting Policies for IPv6 resources.</p>
Core Access - Passthrough proxy	
	<p>The system passthrough proxy modes are based on hostnames or ports, not IP addresses. Therefore, the implementation for IPv6 does not require reconfiguration of the system after upgrade. Note, however, that in virtual hostname mode, your DNS server must be configured to resolve the virtual hostname to the system IP address, which can be an IPv4 or IPv6 address. Update entries in your DNS server accordingly.</p> <p>You cannot configure Passthrough Proxy Policies for IPv6 resources.</p>
Core Access - Hosted Java applets	

Feature	Summary
	<p>The implementation for IPv6 does not require reconfiguration of the system after upgrade. All hosted Java applets, including the premier Java RDP applet, work on IPv4 or IPv6 clients.</p> <p>You cannot configure policies that require access to hosted Java applets at IPv6 addresses.</p>
User role VPN Tunneling options	<p>Route Precedence:</p> <p>If Tunnel Route is selected, the client cannot access its local IPv6 network and IPv6 traffic is blocked, except DHCPv6, ICMPv6, and loopback traffic going to the physical adapter. If Route Monitoring is enabled, only IPv4 route monitoring is performed.</p> <p>If Endpoint Route is selected, the client can access its local IPv6 network. The following options should be disabled for IPv6 resources.</p>
Role/Realm Source IP restrictions	<p>You can specify IPv4 or IPv6 Source IP restrictions at both the role and the realm level.</p> <p>If the device is deployed behind a NAT64 device, it sees traffic coming from an IPv4 address. In this case, your Source IP restrictions should be based on the NATed IPv4 addresses.</p>
Session roaming	<p>You can manage session roaming across IPv6 subnets. If you enable unlimited session roaming, a session is maintained within an IPv4 network, within an IPv6 network, or from IPv4 to IPv6 and vice versa. If you configure limited session roaming, you can specify IPv4 or IPv6 subnets within which the session is maintained. However, with limited session roaming, you cannot allow sessions to roam from IPv4 to IPv6 networks, or vice versa.</p>
Logging	<p>The logging system can process and parse logs containing IPv6 addresses. Ivanti Connect Secure supports communication with external log systems and utilities, such as syslog, SNMP, and archiving that are reachable by IPv4 only.</p>
Network tools	<p>ping6 and traceroute6 were added to the admin graphical user interface console network tools page.</p>

IPv6 Feature Configuration Task Summary

IPv6-related features are not enabled by default. After you upgrade the system software, perform the tasks summarized in The following table describes how to make the device ready for IPv6 traffic.

The following table lists the IPv6 Feature Configuration Task Summary

Action	Documentation
Enable IPv6 for the external port and configure an IPv6 address.	Configuring the External Port
Enable IPv6 for the internal port and configure an IPv6 address.	Configuring the Internal Port
Enable IPv6 for the management port and configure an IPv6 address.	Using the Management Port
Configure IP aliases and IPv6 addresses for virtual ports.	Using Virtual Ports
Re-create a cluster deployment with IPv6 configuration for external interfaces.	Clustering
If you use source IP policies, configure them so that source IP restrictions are based on IPv6 addresses.	Specifying Source IP Access Restrictions
Configure IPv6 address assignment for VPN Tunneling Connection Profiles. DHCPv6 is not supported. Also note that the IP address server configuration on the System > Network > VPN Tunneling page does not support filtering for IPv6 address pools.	VPN Tunneling
If you permit roaming sessions but limit to roaming within the specified subnet, configure the role session option so that the subnet is defined by netmask for IPv4 and prefix length for IPv6 networks.	Specifying Role Session Options
View and manage the neighbor discovery cache. You can view discovered neighbors or clear the entire cache, but you cannot add neighbors or delete individual entries.	Managing the Neighbor Discovery Table
View IPv6 routes in the IP route table. You can view discovered IPv6 routes, but you cannot add or delete them from the route table.	Managing the Routes Table

Action	Documentation
Review logs. The logging infrastructure accommodates IPv6 addresses, and you can create custom filters based on IPv6 address patterns.	Using Log Filters
Become familiar with IPv6 network connectivity test tools, such as ping6 and traceroute6.	Using Network Troubleshooting Commands

Configuring SSL Options

Use the System > Configuration > Security > SSL Options page to change the default security settings. We recommend that you use the default security settings, which provide maximum security, but you may need to modify these settings if your users cannot use certain browsers or access certain Web pages.

TLS_DHE_RSA_WITH_AES_128_CBC_SHA TLS_DHE_RSA_WITH_AES_256_CBC_SHA cipher suites are supported. Both these ciphers use RSA for server authentication and ephemeral Diffie-Hellman (DHE) for key exchange. RSA server certificate is required for these ciphers. Only TLS_DHE_RSA_WITH_AES_256_CBC_SHA is available with the **Accept 168-bit and greater** option. In the Custom SSL Cipher configuration, TLS_DHE_RSA_WITH_AES_128_CBC_SHA is available only when **AES-Medium** is selected and TLS_DHE_RSA_WITH_AES_256_CBC_SHA is available only when **AES-High** is selected. Both ciphers are lower in priority over the other widely used cipher suites.

Enabling Granular Cipher Selection for Setting the Security Options

Granular cipher selection provides an administrator the ability to select specific ciphers and the preferred ordering of the selected ciphers. This feature also provides presets like Suite-B and PFS. There are two tabs, Inbound OpenSSL options and Outbound OpenSSL options. With this feature admins can select the ciphers that TLS/SSL connections will use. The Inbound OpenSSL options apply to all incoming connections. Outbound OpenSSL options apply to the following services:

- Rewriter
- ActiveSync
- SCEP
- Syslog

- LDAPS

i FIPS Mode Settings is common for both Inbound and Outbound SSL Options.

A common cipher library has been added which can be used by both, the inbound and outbound connections. The outbound options are listed in a separate tab next to the inbound settings. The outbound settings have presets for High and Medium ciphers along with custom options. There is no PFS or SuiteB presets on the outbound side. From 8.2R3 release onwards, support for preset Low has been removed and the same can be configured using Custom SSL Cipher Selection option. For the SuiteB preset to work, IVE should have ECC Device Certificate mapped to Internal or External Port. SuiteB preset does not work if the ECC Device Certificate is mapped only to virtual port.

SSL FIPS Mode option


Enabling Inbound SSL Options

Only when FIPS mode is turned on, the FIPS compliant ciphers are available to be chosen from the Supported Ciphers panel. FIPS mode is editable only on the inbound option page.

To set the security options with Inbound SSL Options:

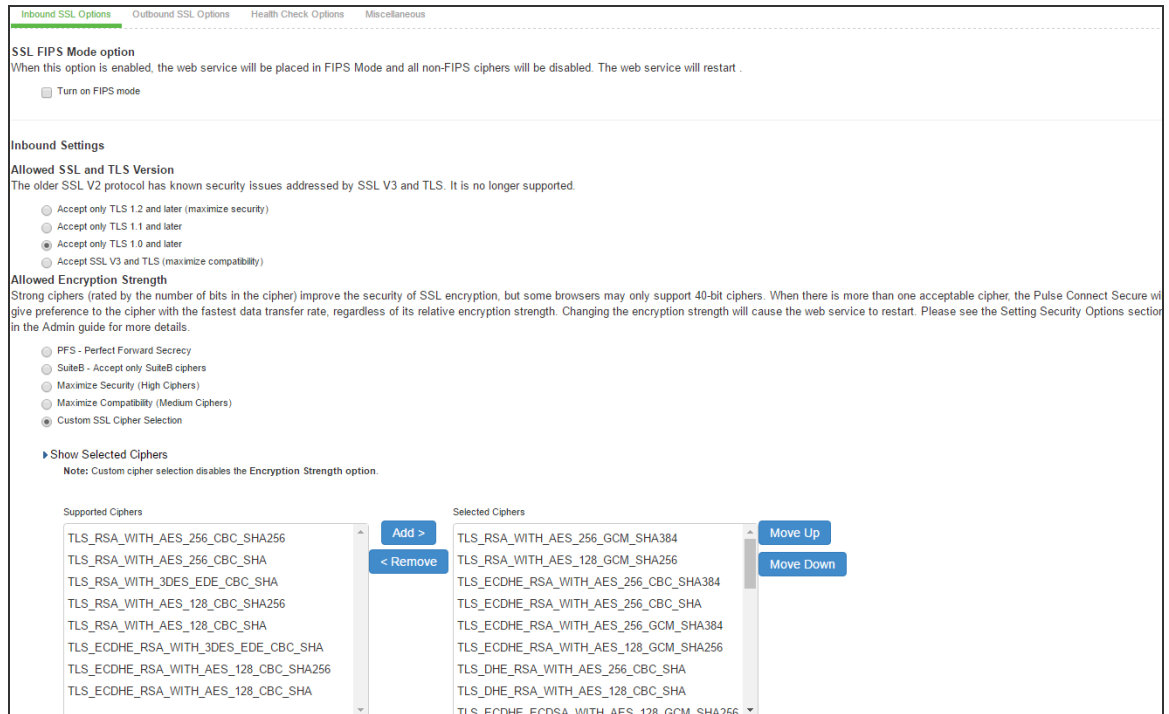
1. In the admin console, select **System > Configuration > Security > Inbound SSL Options**.
2. Under Allowed Encryption Strength choose **Custom SSL Cipher Selection**. See the following figure.

The following figure depicts the Setting Custom SSL Cipher Selections:




3. The two panels of **Supported Ciphers** and **Selected Ciphers** are displayed. Supported ciphers has the entire list of ciphers supported for the selected SSL or TLS version. Selected ciphers list the currently selected ciphers list. The following figure shows the two panels (Supported Ciphers and Selected Ciphers). Note that the Selected Ciphers and Supported Ciphers List will also be displayed for all Preset like PFS or SuiteB or Medium or High.

The following figure depicts the Supported Ciphers and Selected Ciphers Panels:



4. To add a cipher to be used in order to secure a connection, click on the cipher string on the left panel and then click on the **Add >** or double click on the cipher name in the left panel. See the following figure.
5. To remove the cipher, click on the cipher name on the right panel and then click on the **< Remove** button or double click on the cipher name on the right side. See the following figure.
6. The selected ciphers on the right are listed in order of their priority from top to bottom. To change the priority of the ciphers, click on the cipher name and then click on **Move Up** to increase priority or the **Move Down** button to decrease the priority. See the following figure.


The following figure depicts the Setting Custom SSL Cipher Selections:



7. If you are using client certificate authentication (Connect Secure only):

- Select Enable client certificate on the external port under ActiveSync Client Certificate Configuration. See the following figure.
- Move p_ecdsa256 to the Selected Virtual Ports column.

ActiveSync Client Certificate Configuration:



8. Click **Save Changes**.

A list of the custom ciphers to be used on the device's port is displayed in the order the web server will select them. Note that Suite B ciphers are listed on top. See the following table that depicts end users who now log in to external virtual port p_ecdsa256 must have at least one of the listed ciphers installed on their browser or else they cannot log in to the server.

The following figure depicts Confirming Custom Ciphers:

```

Configuration > Security > SSL Options
SSL Options

? Confirm Change to Custom Ciphers
Are you sure you want to use custom ciphers selection? If your browser does not support at least one of the ciphers listed below, you will not be able to continue.

TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA
TLS_RSA_WITH_AES_256_CBC_SHA256
TLS_RSA_WITH_AES_256_CBC_SHA
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA
TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA
TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA
TLS_RSA_WITH_3DES_EDE_CBC_SHA
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA
TLS_RSA_WITH_AES_128_CBC_SHA256
TLS_RSA_WITH_AES_128_CBC_SHA
TLS_RSA_WITH_AES_256_GCM_SHA384
TLS_RSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
TLS_DHE_RSA_WITH_AES_256_CBC_SHA
TLS_DHE_RSA_WITH_AES_128_CBC_SHA

```


9. Click **Change Allowed Encryption Strength**.

- When custom ciphers are selected, there is a possibility that some ciphers are not supported by the web browser. Also, if any of ECDH/ECDSA ciphers are selected, they require ECC certificate to be mapped to the internal/external interface. If ECC certificate is not installed, admin may not be able to log in to the box. The only way to recover from this situation is to connect to the system console and select option 8 to reset the SSL settings from the console menu. Option 8 resets the SSL settings to its default. So, the previously set SSL settings are lost. This is applicable only to Inbound SSL settings.
- Pulse Mobile client does not connect to the ICS device if the ciphers selected in Inbound option are not supported by the mobile client.



Enabling Outbound SSL Options

Only for Outbound SSL Settings, we can configure Non FIPS Ciphers when FIPS is Enabled using Custom Cipher Selection Option. Now, there are options to change different SSL/TLS versions and different encryptions in the Outbound SSL Settings. The following figure shows the Outbound SSL Settings.



The following table lists the SSL Options Configuration Guidelines:

Settings	Guidelines
SSL FIPS Mode option	Enable FIPS mode. See the Connect Secure FIPS Level 1 Feature Guide.
Allowed SSL and TLS Version	Specify encryption requirements for clients. By default, the system requires SSL version 3 and TLS. The system honors this setting for all Web server traffic and all types of clients. You can require users who have older browsers that use SSL version 2 to update their browsers, or you can change this setting to allow SSL version 2, SSL version 3, and TLS.

Settings	Guidelines	
Allowed Encryp-tion Strength	<p>Accept only 128-bit and greater-The default. The system gives preference to RC4 ciphers. You can require users to have this level of encryption strength or change this default to an option compatible with the user base.</p> <p>Accept only 168-bit and greater-The system gives preference to 256-bit AES over 3DES.</p> <p>Accept 40-bit and greater-The system gives preference to RC4 ciphers. Older browsers that predate the change in the U.S. export law in year 2000 that required 40-bit cipher encryption for international export, can still use 40-bit encryption.</p> <p>Custom SSL Cipher Selection-Specify a combination of cipher suites for the incoming connection from the user's browser. If you select the AES/3DES option, the system gives preference to 256-bit AES over 3DES.</p> <p>When using 168-bit encryption, some Web browsers may still show 128-bit encryption (the gold lock on the browser status bar) even though the connection is 168-bit. This is typically a limitation of the browser's capability.</p> <p>If you are using the IC6500 FIPS version, you can choose High, Medium, or Low security cipher suites. AES/3DES High and AES Medium are recommended for FIPS deployment.</p>	
Encryption Strength option	<p>Normally, the allowed encryption strength is enforced after an SSL session is established, so that a user connecting with a disallowed encryption strength receives a Web page describing the problem. Enable this option to prevent a browser with a weak cipher from establishing a connection.</p>	
SSL Handshake Timeout option	<p>Determines how many seconds elapse before the SSL handshake times out. The default is 60 seconds.</p>	

Settings	Guidelines	
SSL Legacy Renegotiation Support option	SSL and Transport Layer Security (TLS) renegotiations can be subjected to man-in-the-middle (MITM) attacks that can lead to abuse. A new TLS extension (defined in RFC 5746) ties renegotiations to the TLS connections they are being performed over to prevent these kinds of attacks. The SSL Legacy Renegotiation Support option is enabled by default and allows renegotiation between clients and servers even if they do not support the new TLS extension. Disable this option to not allow renegotiations between clients and servers that do not support the new TLS extension. A web server restart is required when you change the value of this option.	
ActiveSync Client Certificate Configuration	Use these controls to enforce client certificate requirement for activesync access on the selected ports, including virtual ports. When enabled, all ActiveSync clients must present a client authentication certificate to the system to be able to connect using ActiveSync. Non-ActiveSync access (like web browser-based access to the host, NC, JSAM, PSAM, Ivanti, WTS, IKEv2 and so forth) on the port/interface on which the ActiveSync client certificate is required might not work properly. We recommend you use a separate port or interface exclusively for ActiveSync access and then enable client certificate requirement for the port intended for ActiveSync access.	

SSL NDcPP Mode Option

NDcPP mode can be enabled in the Inbound tab with a check box. This status is also applied over to the Outbound tab. Turning on NDcPP automatically turns on FIPS mode and disables SSL/TLS Version TLS1.0 and below. Also, NDcPP Mode allows to choose only 16 Ciphers under Custom Encryption Strength. Turning on the NDcPP check box selects all the NDcPP ciphers by default on both, the Inbound and Outbound sides.



When the NDcPP Mode is enabled, backend server like Windows 2008 R2 which supports the SSL/TLS Version only till TLS1.0 cannot be connected via Rewriter.

syslog-ng server

- Connection to syslog-ng server does NOT get established, since syslog-ng does not support TLSv1.1 and TLSv1.2.

rsyslog

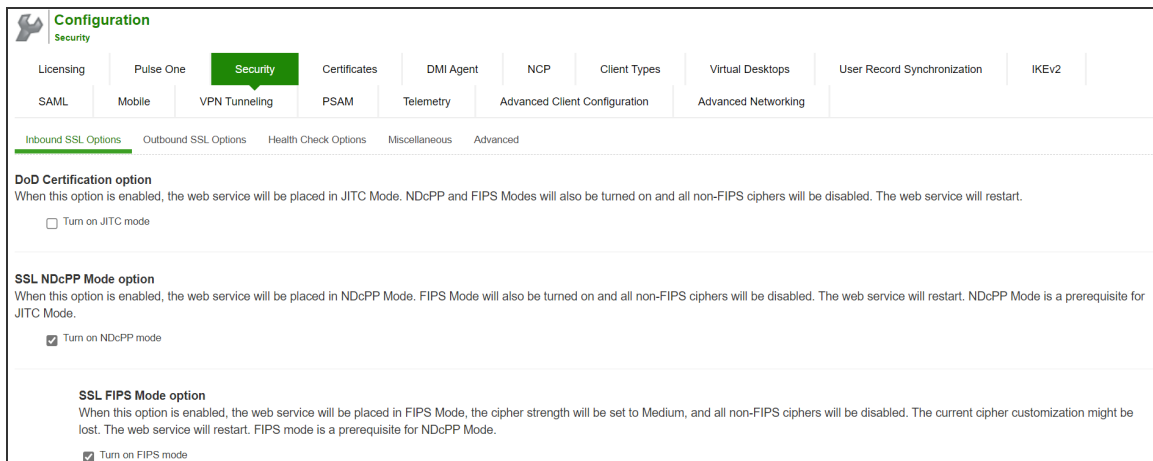
- Supports only till TLSv1.1. So, connection would not get established, if Outbound SSL Options is set to use TLSv1.2.

- To be NDcPP compliant, NTP Update Interval needs to be limited to 60 minutes. This is to avoid the potential drift becoming too excessive.



- For incoming client certificate during client certificate authentication and for incoming server certificate during backend syslog server connection 1024-bit Key Length is not allowed in both NDcPP and FIPS Mode where as SHA1 Signature Algorithm is not allowed only in FIPS Mode and is allowed in NDcPP Mode. This restriction is not applicable for Outgoing Certificates from ICS during SSL Negotiation.

The following figure depicts the SSL NDcPP Mode Option (**System > Configuration > Security > SSL Options**):



Admin Password Storage

NDcPP mandates that admin passwords needs to be scrambled with SHA2 algorithm. So, current SHA1 password scrambling is no longer supported. Password migration is done through double hashing. Existing scrambled passwords stored in the cache are scrambled again with SHA 512.

New passwords will be hashed twice: first with SHA1 and then with SHA512 and then, stored in the cache.

Inbound Settings

When the NDcPP mode is enabled, the following settings appear by default in the Inbound SSL Options page:

- The **Accept only TLS 1.1 and later** is enabled by default in the Allowed SSL and TLS Version settings. Only the **Accept only TLS 1.1** and **Accept only TLS 1.2** options can be chosen. The **Accept only TLS 1.0 and later** and the **Accept SSL V3** and **TLS** (maximize compatibility) are disabled. See the following figure.
- With regards to the Allowed Encryption Strength settings the **Custom SSL Cipher Selection** is enabled by default with NDcPP Ciphers. All other options are disabled.

The following figure depicts the NDcPP Inbound Settings Page (**System > SSL NDcPP Mode Option**):

The screenshot shows the NDcPP Inbound Settings Page. It includes the following sections:

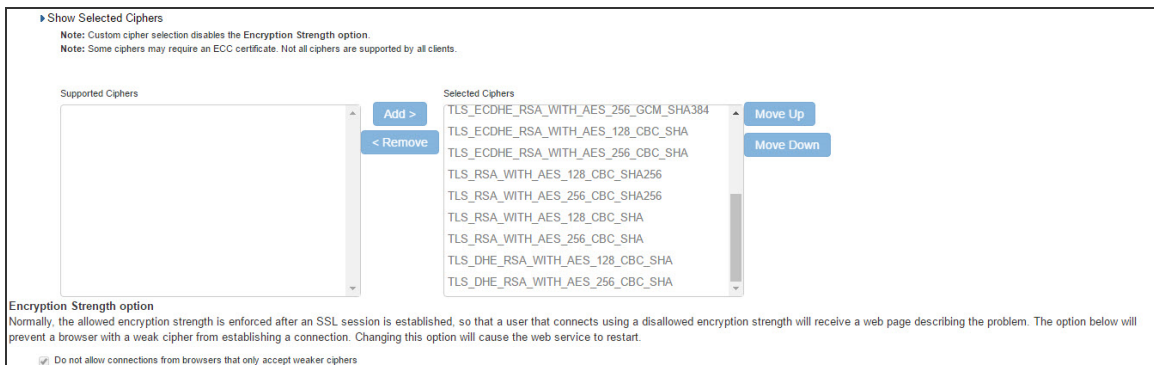
- SSL NDcPP Mode option:** A checkbox labeled "Turn on NDcPP mode" is checked. Below it, a note states: "When this option is enabled, the web service will be placed in NDcPP Mode. FIPS Mode will also be turned on and all non-FIPS ciphers will be disabled. The web service will restart. NDcPP Mode is a prerequisite for JITC Mode."
- SSL FIPS Mode option:** A checkbox labeled "Turn on FIPS mode" is checked. Below it, a note states: "When this option is enabled, the web service will be placed in FIPS Mode, the cipher strength will be set to Medium, and all non-FIPS ciphers will be disabled. The current cipher customization might be lost. The web service will restart. FIPS mode is a prerequisite for NDcPP Mode."
- Inbound Settings:**
 - Allowed SSL and TLS Version:** A note states: "The older SSL V2 protocol has known security issues addressed by SSL V3 and TLS. It is no longer supported." Four radio button options are listed:
 - Accept only TLS 1.2 and later (maximize security)
 - Accept only TLS 1.1 and later** (selected)
 - Accept only TLS 1.0 and later
 - Accept SSL V3 and TLS (maximize compatibility)
 - Allowed Encryption Strength:** A note states: "Strong ciphers (rated by the number of bits in the cipher) improve the security of SSL encryption, but some browsers may only support 40-bit ciphers. When there is more than one acceptable cipher, the Pulse Connect Secure will give preference to the cipher with the fastest data transfer rate, regardless of its relative encryption strength. Changing the encryption strength will cause the web service to restart. Please see the Setting Security Options section in the Admin guide for more details." Five radio button options are listed:
 - FPS - Perfect Forward Secrecy
 - SuiteB - Accept only SuiteB ciphers (Requires an ECC certificate)
 - Maximize Security (High Ciphers)
 - Maximize Compatibility (Medium Ciphers)
 - Custom SSL Cipher Selection** (selected)

The following is a list of Selected Ciphers in the Inbound Settings with the NDcPP mode enabled:

- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA

- TLS_RSA_WITH_AES_128_CBC_SHA256
- TLS_RSA_WITH_AES_256_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384

The following figure depicts the Selected Ciphers in the Inbound Settings with the NDcPP Mode:



Outbound Settings

When the NDcPP mode is enabled, the following settings appear by default in the Outbound SSL Options page:

- The **Accept only TLS 1.1 and later** is enabled by default in the Allowed SSL and TLS Version settings. Only the **Accept only TLS 1.1** and **Accept only TLS 1.2** are editable. The **Accept only TLS 1.0** and later and the **Accept SSL V3 and TLS** (maximize compatibility) are disabled.
- With regards to the Allowed Encryption Strength settings the **Custom SSL Cipher Selection** is enabled by default. All other options are disabled.
- Only the NDcPP ciphers configured in the Outbound SSL options settings are sent in the Outbound connections (ICS -> backend SSL).

The following figure depicts the NDcPP Outbound Settings Page (**System > Configuration > Security > Outbound SSL Options**):

Configuration > Security > Outbound SSL Options

Outbound SSL Options

Configuration
Security

Licensing | Pulse One | **Security** | Certificates | DMI Agent | NCP | Client Types | Virtual Desktops | User Record Synchronization | IKEv2

SAML | Mobile | VPN Tunneling | PSAM | Telemetry | Advanced Client Configuration | Advanced Networking

Inbound SSL Options | **Outbound SSL Options** | Health Check Options | Miscellaneous | Advanced

DoD Certification Mode
JITC is Off

SSL NDcPP Mode
NDcPP is Off

SSL FIPS Mode
FIPS is Off

Outbound Settings

Allowed SSL and TLS Version

The older SSL V2 protocol has known security issues addressed by SSL V3 and TLS. It is no longer supported. These settings apply to the outbound connections that this device makes. No services are restarted with these changes, but the settings will take effect for all new connections established after the change.

Use TLS 1.2 (maximize security)
 Use TLS 1.1 and later
 Use TLS 1.0 and later
 Use SSL V3 and TLS (maximize compatibility)

Allowed Encryption Strength

Please see the Setting Security Options section in the Admin guide for more details.

Maximize Security (High Ciphers)
 Maximize Compatibility (Medium Ciphers)
 Custom SSL Cipher Selection

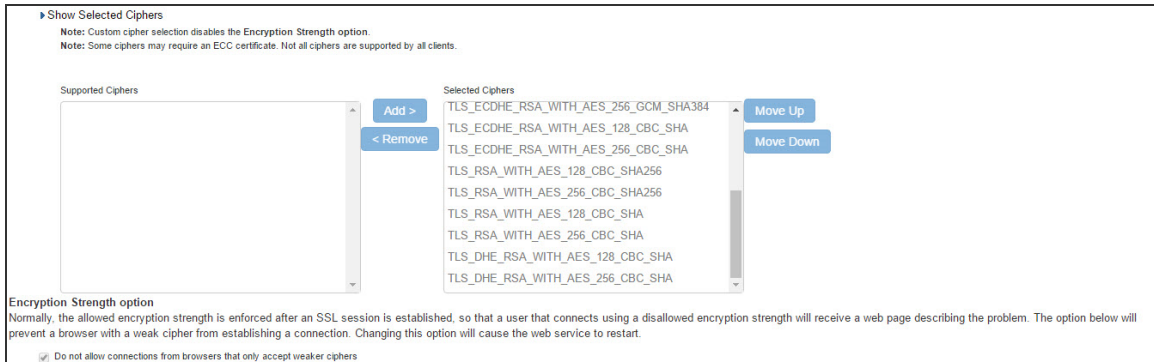
Note that changing any of the above settings might restart some services in the Pulse Connect Secure.

The following is a list of Selected Ciphers in the Outbound Settings with the NDcPP mode enabled:

- TLS_RSA_WITH_AES_128_CBC_SHA
- TLS_RSA_WITH_AES_256_CBC_SHA
- TLS_DHE_RSA_WITH_AES_128_CBC_SHA
- TLS_DHE_RSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
- TLS_RSA_WITH_AES_128_CBC_SHA256
- TLS_RSA_WITH_AES_256_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384

- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384

The following figure depicts the Selected Ciphers in the Outbound Settings with the NDcPP Mode:



Configuring Health Check Options


You can use the System > Configuration > Security > Health Check Options page to configure the following security options:

Enable additional information via healthcheck.cgi-This option is used by entities like load balancers to monitor the health status of the node.

To configure health check options:

1. Select **System > Configuration > Security > Health Check Options** to display the configuration page.

The following figure shows the configuration page:



2. Select the **Enable additional information via healthcheck.cgi** checkbox and **Save Changes**. A URL parameter 'status' needs to be passed to get additional information to the health check url.

For more information about parameters such as CPU usage and number of active sessions use https://<Ivanti Connect Secure>/dana-na/healthcheck/healthcheck.cgi?status=all.

For more information about SBR statistics use https://<Ivanti Connect Secure>/dana-na/healthcheck/healthcheck.cgi?status=sbr

3. Add the relevant IPv4/v6 addresses for which additional information is required to be made available, and click Add.
4. Now click **Save Changes**.

Configuring Miscellaneous Security Options

You can use the System > Configuration > Security > Miscellaneous page to configure the following security options:

- **Persistent cookie options** - You can choose whether to preserve or delete persistent cookies when a session is terminated.
- **Lockout options** - You can configure lockout options to protect the system from denial of service (DoS), distributed denial of service (DDoS), and password-guessing attacks.

- **Last login** - You can choose whether to show users the time and IP address their user ID was used to sign in.
- **X-Frame-Options protection** - You can choose to defend against click-jacking attacks by adding X-Frame-Option header to all the IVE generated pages. If this is not enabled, then only welcome.cgi will have this header.
- **Slow Post Attack Defense** - You can configure to protect against slow-post DOS attacks from non-authenticated users.
- **Host Manifest Integrity Validation** - You can configure to protect against integrity attacks.
- **Integrity checking options** - You can configure scan the system to periodically check for any integrity anomalies. If any anomaly found, information is displayed in the dashboard.

On low-end machines, a reduction in through-put may be seen. To overcome such issues, opt for scheduled scan option in off-peak hours.

To configure cookie and lockout options:

1. Select **System > Configuration > Security > Miscellaneous** to display the configuration page.

The following figure shows the configuration page.

2. Complete the configuration as described in the following table.
3. Save the configuration.

The following figure depicts the Miscellaneous Security Options Configuration Page:

Configuration > Security > Miscellaneous

Miscellaneous



Licensing	Pulse One	Security	Certificates	DMI Agent	NCP	Client Types	Virtual Desktops	User Record Synchronization	IKEv2
SAML	Mobile	VPN Tunneling	PSAM	Telemetry	Advanced Client Configuration	Advanced Networking			
Inbound SSL Options		Outbound SSL Options		Health Check Options		Miscellaneous	Advanced		

Delete all cookies at session termination

For convenience, some persistent cookies (the last realm cookie and the last sign-in URL cookie) are set on the user's machine. If you desire additional security or privacy, you may choose to not set them.

- Delete all cookies at session termination (maximize security)
- Preserve cookies at session termination (maximize usability)

Include Ivanti Connect Secure's session cookie in URL

Depending on privacy settings, Mozilla may withhold cookies from the Ivanti Connect Secure and JVM, thereby preventing users from running java applets and java-enabled applications such as J-SAM. To enable users to launch these applications without changing their browser settings, the Ivanti Connect Secure can include the user's session cookie in the URL that launches java applets and java-enabled applications.

- Include session cookie in URL (maximize compatibility)
- Do not include session cookie in URL (maximize security)

Lockout options

The following settings determine how failed sign-in attempts are handled. When the number of allowed attempts is exceeded, the IP address that is used for signing-in will be temporarily locked to prevent automated sign-in attacks.

Rate: per minute 1-2147483647 Rate of failed attempts

Attempts: 2-2147483647 Initial trigger of failed attempts

Lockout period: minutes (1-10080 minutes)

Last Login options

The following settings determine whether to show the user's last login time and source IP address details on the user's bookmark page. For Admin users this information will be displayed on the System Status page. These settings do not apply to the custom start page option on Role UI options page.

- Show last login time on user's bookmark page
- Show last login IP address on user's bookmark page

X-Frame-Options protection

Enable X-Frame-Options protection to defend against click-jacking attacks by adding X-Frame-Option header to all the IVE generated pages. If this is not enabled then only welcome.cgi will have this header

- Enable X-frame options protection

SYN FLOOD, SMURF, SSL Replay Attack Audit Logs

Enable SYN Flood, Smurf and SSL Replay attack audit logs. Turning this on can have performance and resource impact. Even when turned off, the device is always protected against these attacks. This option controls only the logging for these attacks. This option needs to be on when the device is in JTC Mode

- Enable SYN Flood, SMURF, SSL Replay Attack Audit

Slow post attack defence

The Ivanti Connect Secure is vulnerable to a slow post HTTP attack, which is a kind of Denial-of-Service (DOS) attack in which the attacker slowly sends HTTP requests in small pieces, keeping server resources busy and maxing out concurrent connection pools. The following countermeasures are supported:

1. Set a smaller maximum wait time (Timeout) for an unauthenticated post connection to complete.
2. Set a smaller maximum request size. Unauthenticated requests exceeding set values will be dropped.

NOTE: Very small values for either parameter may result in legitimate requests being dropped. The settings should minimally be slightly higher than lifetime statistical medians.

Timeout: Seconds 3 - 180 Seconds

Maximum request size: Bytes 256 - 524288 Bytes

HSTS

HTTP Strict Transport Security (HSTS) is a HTTP special response header which will prevent any communications over HTTP and also prevents HTTPS click through prompts on browsers.

Max Age: Days 0 - 365 days

- Enable includeSubDomain directive
- Enable preload directive

Host Manifest Integrity Validation

- Enable Host Manifest Integrity Validation to stop booting if manifest integrity validation fails

Host Header Validation

- Enable Host header validation to block open redirect attacks

Username Validation

- Enable Username validation to block unauthorised access

Runtime Integrity Scanner Interval

- Periodic Scan Scheduled Scan
- Every 1 Hour
- Every 2 Hours
- Every 6 Hours
- Every 12 Hours
- Every 24 Hours

Referer Header Validation

- Enable Referer Header validation to block CSRF attacks

Save Changes

The following table lists the Miscellaneous Security Options Configuration Guidelines:

Settings	Guidelines
Delete all cookies at session termination	
Delete / Preserve	For convenience, the system sets persistent cookies on the user's machine to support functions such as multiple sign-in, last associated realm, and the last sign-in URL. For additional security or privacy, you can choose not to set them.
Include Ivanti Connect Secure's session cookie in URL	
Include / Not Include	Mozilla 1.6 and Safari may not pass cookies to the Java Virtual Machine, preventing users from running JSAM and Java applets. To support these browsers, the system can include the user session cookie in the URL that launches JSAM or a Java applet. By default, this option is enabled, but if you have concerns about exposing the cookie in the URL, you can disable this feature.
Lockout options	
Rate	Specify the number of failed sign-in attempts to allow per minute.
Attempts	Specify the maximum number of failed sign-in attempts to allow before triggering the initial lockout. The system determines the maximum initial period of time (in minutes) to allow the failed sign-in attempts to occur by dividing the specified number of attempts by the rate. For example, 180 attempts divided by a rate of 3 results in an initial period of 60 minutes. If 180 or more failed sign-in attempts occur within 60 minutes or less, the system locks out the IP address being used for the failed sign-in attempt.
Lockout period	Specify the length of time (in minutes) the system must lock out the IP address.
Last Login options	
Time / IP Address	Display the day and time and IP address the user last logged in to the system. For users, this information appears on their bookmark page. For administrators, this information appears on the System Status Overview page. These settings do not apply to the custom start page option on Role UI Options page.
X-Frame-Options protection	

Settings	Guidelines
Enable X-Frame-Options protection	By default, the Enable X-Frame-Options is checked. If the admin does not want to have this protection, they can uncheck this option. The X-Frame-Options HTTP response header can be used to indicate whether or not a browser should be allowed to render a page in a <frame>, <iframe> or <object>.
Slow Post Attack Defence	
Timeout	By default, the POST body is received within 10 seconds. If the browser is unable to send the POST body within 10 seconds the connection is eventually dropped. (Configurable from 3 - 60Sec)
Maximum Request Size	By default, now a connection is directly rejected if it tries to POST more than 4KB in POST body (Configurable from 256 Bytes to 24 KB)
HSTS	
Max Age	Specify the maximum age for HSTS. It can be disabled by configuring max age as 0.
Enable includeSub-domain directive	Select the check box to enable/disable the includeSubdomain directive. By default, it is turned off.
Enable preload directive	Select the check box to enable/disable the preload directive. By default, it is turned off.
Host Manifest Integrity Validation	
Enable Host Manifest Integrity Validation to stop booting if manifest integrity validation fails	Select the check box to enforce host manifest integrity validation. By default, it is turned off. The following integrity checks are performed: Checks the SHA512 digital signature of the manifest file. Checks the SHA256 digest of each individual file entries in the manifest. If enabled and integrity check fails, admin needs to roll back to previous working package or perform factory reset.
Host Header Validation	
Enable Host header validation to block open redirect attacks	Select the check box to enforce host header validation. By default, it is turned off.

Settings	Guidelines
	When Host header validation is enabled, every http request will be validated against hostnames and IP v4/v6 addresses known to the ICS server. If match is not found, the request will be dropped and logs are recorded in admin access logs and user access logs, and a response will be sent back to client.
Username Validation	
Enable Username validation to block unauthorised access	Select the check box to enforce username validation for usage of unsupported characters. Max allowed length for username is 128 characters.
Runtime Integrity Scanner Interval	
Periodic Scan	Select the time interval to run the integrity scanner during run time. For example: Select 2 hrs to run the integrity scanner every 2 hrs.
Scheduled Scan	Select to run integrity scanner at a specified time everyday. For example: When 13 hr 25 min is specified, the scanner runs at the same time everyday.
Referer Header Validation	
Enable Referer Header validation to block CSRF attacks	Select the check box to enable referer header validation.

Scenario Illustrating Lockout Settings Workflow

The following scenario illustrates how lockout settings work. For example, assume the following settings:

- Rate = 3 failed sign-in attempts per minute
- Attempts = 180 maximum allowed in initial period of 60 minutes (180/3)
- Lockout period = 2 minutes

The following sequence illustrates the effect of these settings:

1. During a period of 3 minutes, 180 failed sign-in attempts occur from the same IP address. Because the specified value for Attempts occurs in less than the allowed initial period of 60 minutes (180/3), the system locks out the IP address for 2 minutes (fourth and fifth minutes).
2. In the sixth minute, the system removes the lock on the IP address and begins maintaining the rate of 3 failed sign-in attempts/minute. In the sixth and seventh minutes, the number of failed sign-in attempts is 2 per minute, so the system does not lock the IP address. However, when the number of failed sign-in attempts increases to 5 in the eighth minute, which is a total of 9 failed sign-in attempts within 3 minutes, the system locks out the IP address for 2 minutes again (ninth and tenth minutes).
3. In the eleventh minute, the system removes the lock on the IP address and begins maintaining the rate of 3 failed sign-in attempts per minute again. When the rate remains below an average of 3 per minute for 60 minutes, the system returns to its initial monitoring state.

Configuring Custom HTTP Headers

Ivanti Connect Secure supports several HTTP headers, which are sent in response to the client request. There are several more headers built to improve security and prevent attacks like XSS. The Custom HTTP Headers configuration enables the administrator to add new headers that they want to enforce.

To configure custom HTTP header:

1. Select **System > Configuration > Security > Advanced**.

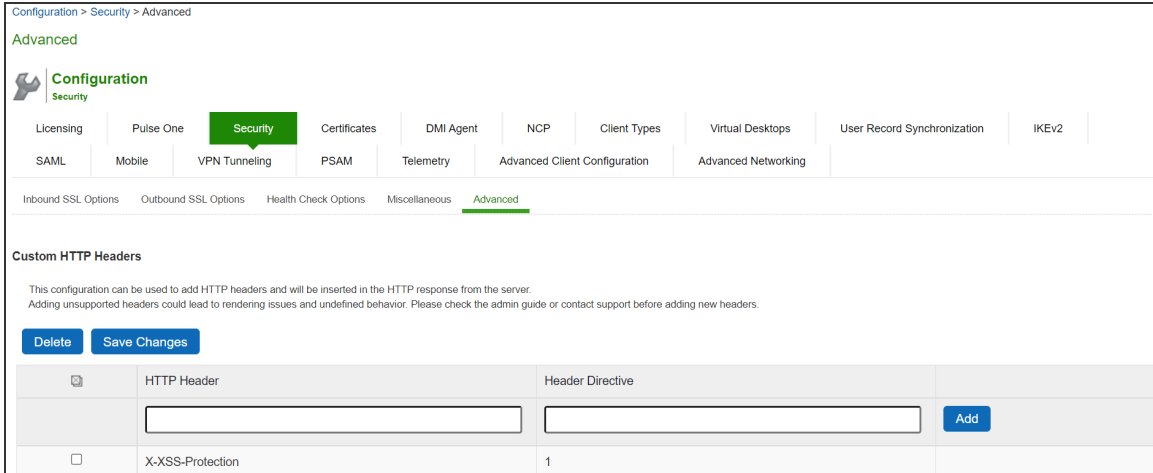
The following figure shows the configuration page.

2. In the Custom HTTP Headers section, enter the HTTP header name and the directives along with the values.
3. Click **Add**.
4. Multiple headers can be added or removed. After adding the headers, click **Save Changes**.



- Administrator should ensure the correctness of the values that they enter, as the system validation on the input values is limited.
 - If the administrator configured HTTP header seems to affect the way the page is rendered or is locked out, use the console option to reset the custom HTTP header values.
-

The following figure depicts the Custom HTTP Headers Page:



The following table lists the OWASP recommended headers:

Header	Supported Browsers
HPKP	Firefox, Chrome, Opera
X-XSS-Protection	Chrome and IE
X-Content-Type-Options	Firefox, Chrome, Opera and IE
Content-Security-Policy	All major browsers
X-Permitted-Cross-Domain-Policies	Not supported
Referrer-Policy	Chrome, Firefox and Opera
Expect-CT	Chrome and Opera
Feature-Policy	Not supported
HSTS	
X-Frame-Options	

Configuring NCP and JCP

The following types of internal protocols are used to communicate between Ivanti Connect Secure and client applications:

- Network Communications Protocol (NCP)-Standard NCP has been replaced with oNCP. Windows client applications, PSAM, and Terminal Services fallback to NCP if oNCP fails.
- Optimized NCP (oNCP)-Optimized NCP (oNCP) significantly improves the throughput performance of the client applications over NCP because it contains improvements to protocol efficiency, connection handling, and data compression. Windows client applications, PSAM, and Terminal Services use oNCP by default.
- Java Communications Protocol (JCP)-JCP is the Java implementation of standard NCP. The system uses JCP to communicate with Java client applications, JSAM, and the Java Content Intermediation Engine.

To set NCP options:

1. In the admin console, choose **System > Configuration > NCP**.
2. (Windows clients) Under NCP Auto-Select, select:
 - **Auto-select Enabled** (recommended)-Use the oNCP by default. If you select this option, the system uses oNCP for most client/server communications and then switches to standard NCP when necessary. The system reverts to NCP if the user is running an unsupported operating system, browser type, or combination thereof, or if the client application fails to open a direct TCP connection to the device for any reason (for instance, the presence of a proxy, timeout, disconnect).
 - **Auto-select Disabled**-Always use standard NCP. This option is primarily provided for backwards compatibility.



If you are using Network Connect to provide client access, we recommend that you exercise caution when employing the Auto-select Disabled option, as Mac and Linux clients cannot connect using the traditional NCP protocol. If you disable the oNCP/NCP auto-selection feature and a UDP-to oNCP/NCP fail-over occurs, the system disconnects Macintosh and Linux clients because it fails over from UDP to NCP (instead of oNCP), which does not support these users.

3. (Java clients) Under Read Connection Timeout, set the timeout interval for Java clients (15-120 seconds). If client-side secure access methods do not receive data from the system for the specified interval, they try to reestablish a connection. Note that this value does not apply to user inactivity in client applications.
4. (Windows clients) Under Idle Connection Timeout, set the idle connection interval. This timeout interval determines how long the system maintains idle connections for client-side Windows secure access methods.
5. Click **Save Changes**.

Using the User Record Synchronization Feature

This topic describes the user record synchronization feature.

User Record Synchronization Overview

The user record synchronization feature promotes a more consistent user experience by allowing users to retain their bookmarks and individual preferences regardless of which device they log in to.

User record synchronization relies on client-server pairings. The client is the Connect Secure device that users log in to start their remote access. Each client is associated with one primary server and one backup server to store user record data. Clients can be individual appliances or a node within a cluster.

A server in this instance is the Connect secure device that stores the user data records. Each server can be configured to replicate its user record data to one or more peer servers. Servers are identified by a user-defined logical name. The same logical name can be assigned to more than one authentication server to let you associate authentication servers of different types to the same user. For example, SA1 is an ACE authentication server with user1 who creates a bookmark to www.ivanti.com. SA2 is an Active Directory authentication server with the same user1. For the www.ivanti.com book-mark to be transferred from SA1/ACE/user1 to SA2/AD/user1 you would assign the logical name "Logi-cal1" to both the ACE server on SA1 and the Active Directory server on SA2.



Cluster VIPs cannot be used as the IP for synchronizing between clients and peers servers.

As long as the logical name is the same, the authentication servers can be different types and different server names and still be associated with a common user. The username must be the same for user record data to be synchronized across the servers. The logical authentication server (LAS) and username combination is what uniquely identifies a user record.

The following user records are synchronized between the client and server:

- Bookmarks
 - Web
 - File
 - Terminal Services
 - JSAM
- Preferences
- Persistent cookies
- Cached passwords

User session data is not synchronized. Persistent cookies, if changed, are synchronized when the user session terminates. All other modifications to the user records are synchronized immediately. User records are stored in cache on the client node prior to being pushed to the servers.

When a user logs in to a client, their data is pulled from the associated server. The pull is performed in the background and does not delay the login process. Users using browsers that do not support JavaScript must manually refresh the index page for updated bookmarks and preferences to appear. For browsers that support JavaScript, users may see a spinning progress indicator and their home page will refresh automatically with updated bookmarks and preferences.

Clients and servers need not be installed with the same system software version.



User record synchronization uses port 17425. This port number is not configurable. If you are deploying across a firewall, configure your firewall to allow traffic on this port.

To set up user record synchronization, you perform the following tasks:

1. Enable user record synchronization for each participating client and server, identify which ones are the client and which ones are the server and assign a node name to each client and server.
2. Create a shared secret that is used to authenticate the client with the server and the server to its peer servers.
3. On each server, define which clients and peers are allowed to communicate with the server.
4. On each client, define the servers that handle records for each LAS server.

When enabling this feature, you have several options to initialize the user record database. You can:

- populate the database using user records located in the cache of the client systems.

- populate the database use user records located in the cache of the server systems.
- don't pre-populate the database but populate it as users log in and out of the client system.

If you choose the last option, users may not be able to view their saved bookmarks and preferences until the next time they log in, depending on which client they log in to.



User records may not synchronize if the time clocks on the devices are not in sync. We recommend that you use the same NTP server for each node participating in user record synchronization to keep times accurately adjusted.

The user record synchronization feature will not start automatically after importing a system configuration that has this feature enabled. The workaround is to disable user record synchronization and then enable user record synchronization from the user interface after the configuration import.

Configuring the User Record Synchronization Authentication Server

To set up the authentication server you must define its logical name:

1. Select **Authentication > Auth Servers**.
2. Click the name of the authentication server you want assign a LAS name.
3. By assigning the authentication server a LAS name, all users that authenticate using the authentication server are associated with this LAS. In this instance, we are referring to the client nodes, not the user record synchronization server nodes.
4. Select the **User Record Synchronization** check box.
5. Enter a logical name to identify this server.

This allows you to share user record data across authentication servers on different Connect Secure devices. By assigning a LAS name to an authentication server, you are implicitly assigning it to all users that authenticate with that auth server. The combination of the user's login name and their LAS name uniquely identifies the user's user record across all user record synchronization servers.

6. Click **Save Changes**.

Configuring the User Record Synchronization Server

To set up the user record synchronization server you must define its peer nodes (optional) and the clients that can access this server.

1. Select **System > Configuration > User Record Synchronization > This Server**.
2. Enter the peer server's node name and IP address, then click **Add**. To specify more than one peer server, enter each server's node name and IP address individually and click Add. There is no limit on the number of peer servers you can add.

Data is replicated from the primary or backup server to its peer servers. If the primary is not available, user data is sent to the backup. User data is then replicated to the peer servers.

3. For each client you want synchronized with this server, enter the client's name and IP address and click Add.

Once added, peer servers will have a colored icon next to their name indicating their connection status. Node status is provided to client nodes and LAS mapping servers as well.

Color	Description
Green	Connected
Yellow	Connecting
Gray	Not connected

Configuring the User Record Synchronization Client

To set up the client, you select the primary and backup server you want this client to synchronize with:

1. Select **System > Configuration > User Record Synchronization > This Client**.
2. Select the LAS name you want to synchronize and enter the primary IP of the user record server that will serve the user records. If you prefer to synchronize with any available server, select **Any LAS**.
3. Enter the primary and optionally a backup server's IP address and then click **Add**.

Even if you select Any LAS, you must enter a primary server IP address.

Once added, the primary and backup servers have a colored icon next to their name indicating their connection status.

Configuring the User Record Synchronization Database

With the Database tab, you can delete inactive records from the client cache, retrieve statistics about the database, export and import the data and remove user data from the server's database.

To configure the database:

1. Select **System > Configuration > User Record Synchronization > Database**.
2. Select **Auto-delete inactive synchronized user records from the Cache** to remove inactive user records from the cache. This option does not remove user records from the user record database.

When this option is selected, the system performs a check every 15 minutes and deletes user records that meet all of the following criteria:

- There are no active user sessions associated with the user record.
 - The user record does not have any custom settings, or the latest version of the user record has been synchronized with the user record database.
 - The authentication server associated with the user record database does not have type "local". For example, the "System Local" auth server that is part of the default configuration has a "local" type, so any user records associated with that auth server will not be auto-deleted. However, user records associated with external authentication servers like Radius or LDAP may be deleted, depending on the two prior criteria.
3. Select **Auto-delete user records from the local synchronization database that have been idle for X days** to permanently remove user records from the database located on the server. Enter the number of days user records must be inactive before being deleted.

In this instance, "inactive" means that no client has pulled the user record or pushed any modifications to the user record in X days.

4. Click **Retrieve Statistics** to display the number of records in the database. You cannot edit or view records in the database.
5. Under Export, you export user records to a file. The user records can be exported from the user record database, or from the cache. The exported file can be used to pre-populate the user record database on another node.
 - Enter the LAS name of the user records you want to export. If you leave this field blank, all user records are exported. If you enter a LAS name, only user records with the entered LAS name are exported.
 - To encrypt the exported data, select the **Encrypt the exported data with password** checkbox and enter the password.
 - Click **Export** to export the user records from the specified source (cache or database). You will be prompted where to save the file.

6. Under Import, you import user records into the synchronization database. The user records can be imported from a file or from the cache. Use the Import operation to pre-populate the user record database with user records exported from another node, or with user records from the cache.

- Click **Browse** to locate the exported file and enter the password if the exported file was encrypted with a password.
- Select the **Override Logical Auth Servers in imported user records** with check box to replace the LAS name in each imported user record with the LAS name entered.

For example, you change the LAS name, use this option to update the user records with the new name.

- Click **Import**.
7. Under Delete, specify which user records to permanently remove from the user record database. The options you select apply only to the user record database associated with this server:
 1. Select **User record with login name and Logical Auth Server** to remove a specific record. The login name and LAS name together uniquely identify a user record. Select this option to remove that record (if it exists).
 2. Select **User records with Logical Auth Server** to delete all user records with the specified LAS name.
 3. Select **All user records** to permanently remove user records from the database on this node.
 4. Click **Delete**.

Enabling User Record Synchronization

The first step in enabling user record synchronizing is to define the node name and the shared secret used to authenticate between the clients and the servers:

1. Select **System > Configuration > User Record Synchronization > General**. See the figure underneath in this section.
2. Select the **Enable User Record Synchronization** check box.
3. Enter a unique node name. This name is used when associating a client with a server and is different from the logical name assigned to a server. This node name is also not the same as the cluster node name.

4. Enter the shared secret and confirm it.

The shared secret is the password used to authenticate the client with its servers and the primary server with its peer servers. Use the same shared secret for all clients and servers participating in user record synchronization.

5. Select whether this node is client only or if this node acts as both a client and server.
6. Click **Save Changes**.



If you need to make any changes in this window at a later time, you must clear the Enable User Record Synchronization check box and click Save Changes. Make your edits, select the Enable User Record Synchronization check box and save your changes.

Once you enter a name and shared secret, you cannot clear these fields.

The following figure depicts the User Record Synchronization General Settings Configuration Page:

Scheduling User Record Synchronization Backup

You can configure periodic backups of the user record database. User record synchronization backup can be enabled only on a user record synchronization server.

To back up the user record database:

1. Ensure the system is set up as a user record synchronization server. See [System > Configuration > User Record Synchronization](#).
2. Select **Maintenance > Archiving > Archiving Servers**.
3. Select the **Archive User Record Synchronization Database** check box.

4. Specify an archive schedule. Through the options, schedule archives on any combination of weekdays including weekends.



If you schedule an archival operation to occur during the hour that your system switches to Daylight Savings Time (DST) the operation may not occur as scheduled. For example, if your system is set to change to DST at 1:00 a.m. and you have scheduled an archival operation to occur at any time between 1:01 a.m. and 1:59 a.m., the operation is not accomplished, because at 1:00 a.m. the system clock is moved forward to 2:00 a.m. and the system never reaches your archival time for that date.

5. Define a specific time when you want the system to archive data or elect to archive data every hour, which produces twenty-four files with unique timestamps.



We recommend you schedule an archival operation during hours when traffic is light in order to minimize its impact to your users. The automatic archiving process compresses files and, if the system is busy, can degrade performance for users. Also, a cluster node may appear unresponsive if the system is busy with traffic and performing archiving simultaneously.

6. Provide a password if you want to encrypt system configuration or user account archives with a pass-word (optional).
7. Click **Save Changes**.

Using IKEv2 Security

This topic describes how to implement IKEv2 security.

IKEv2 Support Overview

This topic gives an overview of support for IKEv2 security.

Understanding IKEv2

IKE or IKEv2 (Internet Key Exchange) is the protocol used to set up a security association in the IPsec protocol suite. Microsoft Windows 7 fully supports the IKEv2 standard through Microsoft's Agile VPN functionality and can operate with a VPN gateway using these protocols. Information on IKE and IKEv2 is widely available on the Internet. It is not the intent of this guide to describe details about IKE and IKEv2.

The system supports IKEv2, enabling interoperability with clients or devices, such as smartphones, that have a standards-based IPsec VPN client.

IKEv2 clients count toward the total number of sessions. Thus, the total number of sessions = number of IKEv2 sessions + number of NCP sessions.

The system supports the following methods for authenticating IKEv2 clients:

- Machine certificate-based authentication
- Authentication using EAP methods



IKEv2 uses port 500 exclusively. Do not configure port 500 in your VPN Tunneling pro-files.

Extensible Authentication Protocol

EAP (Extensible Authentication Protocol) is an authentication framework frequently used in wireless communication. It provides functions and negotiation of authentication methods called EAP methods. Connect Secure supports the following EAP methods:

- EAP-MSCHAP-V2 (Microsoft Challenge-Handshake Authentication Protocol version 2)- a mutual authentication method that supports password-based user or computer authentication. During the EAP-MS-CHAP v2 authentication process, both the client and the authentication server must prove that they have knowledge of the user's password for authentication to succeed. Mutual authentication is provided by including an authenticator packet returned to the client after a successful server authentication. In Connect Secure, the local authentication server and the Active Directory server support EAP-MSCHAP-V2.
- EAP-MD5-Challenge - described in RFC 2284, enables an authentication server to authenticate a connection request by verifying an MD5 hash of a user's password. The server sends the client a random challenge value, and the client proves its identity by hashing the challenge and its password with MD5. EAP-MD5-Challenge is typically used on trusted networks where risk of packet sniffing or active attack are fairly low. Because of significant security vulnerabilities, EAP-MD5-Challenge is not usually used on public networks or wireless networks, because third parties can capture packets and apply dictionary attacks to identify password hashes. Because EAP-MD5-Challenge does not provide server authentication, it is vulnerable to spoofing (a third-party advertising itself as an access point).

Only the local authentication server is supported with EAP-MD5-Challenge.

IKEv2 provides a tunnel mechanism for EAP authentication; it does not perform authentication itself. Instead it proxies EAP messages from a client to the EAP server and back.

- EAP-TLS (Transport Layer Security) - a mutual authentication method that supports certificate-based authentication. EAP-Transport Layer Security Uses the handshake protocol in TLS. During the EAP-TLS authentication process, both client and the authentication server authenticate each other using digital certificates. client generates a pre-master secret key by encrypting a random number with the authentication server's public key and sends it to the authentication server. Both client and authentication server use the pre-master secret key to generate the same master secret key. EAP-TLS is considered to be one of the most secure EAP standards available. The requirement for a client-side certificate is what gives EAP-TLS its authentication strength.

Machine Certificate-Based Authentication

The system supports IKEv2 authentication using machine certificates. Note that only certificate authentication server on Connect Secure supports machine certificate authentication of IKEv2 clients. When using machine certificates for authentication, it is not necessary to configure the Realm/Protocol Set Mapping section under System > Configuration > IKEv2.

Client Requirements

Your IKEv2 client should support the following requirements to work with Connect Secure:

- Ability to establish IPsec Security Associations in Tunnel mode (RFC 4301).
- Ability to utilize the AES 128-bit encryption function (RFC 3602).
- Ability to utilize the SHA-1 hashing function (RFC 2404).
- Ability to utilize Diffie-Hellman Perfect Forward Secrecy in "Group 2" mode (RFC 2409).
- Ability to utilize IPsec Dead Peer Detection (RFC 3706).
- Ability to utilize the MD5 hashing function (RFC 1321).
- Ability to handle Internal Address on a Remote Network utilizing CFG_REQUEST-CFG_REPLY exchange.

Optional but recommended requirements include:

- Ability to adjust the Maximum Segment Size of TCP packets entering the VPN tunnel (RFC 4459).
- Ability to reset the "Don't Fragment" flag on packets (RFC 791).
- Ability to fragment IP packets prior to encryption (RFC 4459).

In addition, your client must support certificate authentication and ESP/SHA1.

Supported Features

The following features are unavailable to the end user since you are using a third-party client that are neither controlled nor configured by Ivanti.

- Host Checker
- Idle timeout notifications
- Upload Logs
- Route monitoring feature of split tunnel
- Windows interactive user logon options
- Session startup scripts
- NCP tunnel mode
- DNS search order
- Proxy server settings

The following table outlines the behavior of the Network Connect client and the IKEv2 client for certain split tunnel options.

The table lists the Split Tunnel Operations with IKEv2 and Network Connect Clients:

Option	IKEv2 Client	Network Connect Client
Disable split tunnel mode	Resource-through tunnel Internet-through tunnel local subnet (client)-through physical adapter	Internet-through tunnel local subnet (client)-through tunnel Resource-through tunnel
Enable split tunnel mode	Resource—through tunnel Internet—through tunnel but fails because the resource is not in split tunnel configuration. local subnet (client)—through physical adapt	Internet—through physical adapter local subnet (client)— through physical adapter

Option	IKEv2 Client	Network Connect Client
Allow local access subnet	Resource—through tunnel Internet—through tunnel local subnet (client)- through physical adapter (same as disable split tunnel mode)	Internet & other traffic—through tunnel local subnet (client)—through physical adapter
Enable split tunnel mode with route monitor (NC proprietary)	Resource—through tunnel Internet—through tunnel but fails because the resource is not in split tunnel configuration. local subnet (client)— through physical adapter Note: route table delete is not monitored.	Resource—through tunnel Internet—through physical adapter local subnet (client)— through physical adapter route table delete is monitored
Enable ST with Allow local access subnet	Resource—through tunnel Internet —through tunnel but fails because the resource is not in split tunnel configuration. local subnet (client)— through physical adapter	Resource—through tunnel Internet—through physical adapter local subnet (client)—through physical adapter

The table below explains the limitations and supported configurations for different IKEv2 clients to work with ICS configured for different IKEv2 authentication:

The following table lists the Limitations and Supported Configurations for Different IKEv2 Clients:

Comparison Parameter	Windows Desktop/Laptop			Windows Mobile Phone		Linux Client	iOS Client	MAC OS Client
	Windows 10-Native Client	Windows 8.1-Native Client	Windows 7-Native Client	Windows 10 Mobile	Windows 8.1 Mobile	Strongswan 5.4.0	iOS 9.X and above	macOS Sierra version 10.12
Client Version	Windows 10-Native Client	Windows 8.1-Native Client	Windows 7-Native Client	Windows 10 Mobile	Windows 8.1 Mobile	Strongswan 5.4.0	iOS 9.X and above	macOS Sierra version 10.12

Comparison Parameter	Windows Desktop/Laptop			Windows Mobile Phone		Linux Client	iOS Client	MAC OS Client
AES128/SHA1 Data Encryption Configuration	Supports only Optional Encryption (connect even if no encryption) Configuration	Supports only Optional Encryption (connect even if no encryption) Configuration	Supports only Optional Encryption (connect even if no encryption) Configuration	Supported	Not Supported	Supported	Supported (this can be configured in the child SA Params in the profile).	Supported
AES256/SHA1 Data Encryption Configuration	Supports all 3 Data Encryption Configuration Optional Encryption (connect even if no encryption)	Supports all 3 Data Encryption Configuration Optional Encryption (connect even if no encryption)	Supports 2 Data Encryption Configuration Optional Encryption (connect even if no encryption) •	Supported	Supported	Supported	Supported (this can be configured in the child SA Params in the profile).	Not Supported

Comparison Parameter	Windows Desktop/Laptop			Windows Mobile Phone		Linux Client	iOS Client	MAC OS Client
	Require Encryption (disconnect if server declines) • Maximum Strength Encryption (disconnect if server declines)	Require Encryption (disconnect if server declines) • Maximum Strength Encryption (disconnect if server declines)	Maximum Strength Encryption (disconnect if server declines) Here Optional Encryption (connect even if no encryption) is not Supported					

Comparison Parameter	Windows Desktop/Laptop			Windows Mobile Phone		Linux Client	iOS Client	MAC OS Client
AES256/SHA256 Data Encryption Configuration	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Supported	Supported	Supported
CA or CA Chain	Need to Import ICS Device Certificate CA in Trusted Root Certificate Authorities Under Computer Account Certificate Store.	Need to Import ICS Device Certificate CA in Trusted Root Certificate Authorities Under Computer Account Certificate Store.	Need to Import • ICS Device Certificate CA in Trusted Root Certificate Authorities Under Computer Account Certificate Store. •	Need to Import • ICS Device Certificate CA in Trusted Root Certificate Authorities Under Computer Account Certificate Store. •	Need to Import • ICS Device Certificate CA in Trusted Root Certificate Authorities Under Computer Account Certificate Store. •	Need to Import ICS Device Certificate CA and SubCA (s) should be placed in cacert directory in pem or cer format.	Need to Import ICS Device Certificate CA / Sub CA (s) Should be Installed through a profile.	Need to Import ICS Device Certificate CA and SubCA (s) should be placed in System > Certificates Key Chain

Comparison Parameter	Windows Desktop/Laptop			Windows Mobile Phone		Linux Client	iOS Client	MAC OS Client
	ICS Device Certificate Sub CA(s) Should be Imported in Intermediate Certificate Authorities Under Computer Account Certificate Store.	ICS Device Certificate Sub CA(s) Should be Imported in Intermediate Certificate Authorities Under Computer Account Certificate Store.	ICS Device Certificate Sub CA(s) Should be Imported in Intermediate Certificate Authorities Under Computer Account Certificate Store.	ICS Device Certificate Sub CA(s) Should be Imported in Intermediate Certificate Authorities Under Computer Account Certificate Store.	ICS Device Certificate Sub CA(s) Should be Imported in Intermediate Certificate Authorities Under Computer Account Certificate Store.			
Client Version	Windows 10-Native Client	Windows 8.1-Native Client	Windows 7-Native Client	Windows 10 Mobile	Windows 8.1 Mobile	StrongSwan 5.4.0	iOS 9.X and above	macOS Sierra version 10.12

Comparison Parameter	Windows Desktop/Laptop			Windows Mobile Phone		Linux Client	iOS Client	MAC OS Client
	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or Encrypting File System (1.3.6.1.4.1.311.10.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or Encrypting File System (1.3.6.1.4.1.311.10.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or Encrypting File System (1.3.6.1.4.1.311.10.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or Encrypting File System (1.3.6.1.4.1.311.10.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or Encrypting File System (1.3.6.1.4.1.311.10.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or Encrypting File System (1.3.6.1.4.1.311.10.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)
Certificate EKU Extension for EAP-TLS	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or Encrypting File System (1.3.6.1.4.1.311.10.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or Encrypting File System (1.3.6.1.4.1.311.10.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or Encrypting File System (1.3.6.1.4.1.311.10.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or Encrypting File System (1.3.6.1.4.1.311.10.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or Encrypting File System (1.3.6.1.4.1.311.10.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)	Client Certificate should have clientAuth (1.3.6.1.5.5.7.3.2)EKU Extension mandatorily. Optionally Certificates can have Secure Email (1.3.6.1.5.5.7.3.4) or serverAuth (1.3.6.1.5.5.7.3.1)

Comparison Parameter	Windows Desktop/Laptop			Windows Mobile Phone		Linux Client	iOS Client	MAC OS Client
	Supported	Not Supported	Not Supported	Supported	Not Supported			
NDcPP Mode	Supported	Not Supported	Not Supported	Supported	Not Supported	Supported	Supported	Not Supported
TLS Version	Supports SSLv3, TLS1.0, TLS1.1 and TLS1.2	Supports only SSLv3 and TLS1.0	Supports only SSLv3 and TLS1.0	Supports SSLv3, TLS1.0, TLS1.1 and TLS1.2	Supports only SSLv3 and TLS1.0	Supports SSLv3, TLS1.0, TLS1.1 and TLS1.2	Supports SSLv3, TLS1.0, TLS1.1 and TLS1.2	Supports only TLS1.0
Suite B Encryption	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
ICS Configured for ECC Device Certificate	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported

Comparison Parameter	Windows Desktop/Laptop			Windows Mobile Phone		Linux Client	iOS Client	MAC OS Client
AES256/MD5 and AES128/MD5 ESP Encryption	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
Client Proxy	Not Working	Working	Working	Not Working	Working	Not Tested	Not Tested	Not Tested
ICS Configured for RSA SHA2 Device Certificate	Supported	Supported	Not Supported	Supported	Supported	Supported	Supported	Supported
ICS Split Tunnel Configuration	Not Supported	Supported	Supported	Not Supported	Supported	Not Tested	Supported	Supported

Comparison Parameter	Windows Desktop/Laptop			Windows Mobile Phone		Linux Client	iOS Client	MAC OS Client
	ICS Configured for Secondary Authentication	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
ICS configured for two or more role mapping roles with "User must select from among assigned roles" option	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not supported

Comparison Parameter	Windows Desktop/Laptop			Windows Mobile Phone		Linux Client	iOS Client	MAC OS Client
	ICS IKEv2 EAP-TLS Configuration	Supports both Machine Certificate Authentication and EAP-TLS Authentication	Supports both Machine Certificate Authentication and EAP-TLS Authentication	Supports both Machine Certificate Authentication and EAP-TLS Authentication	Supports both Machine Certificate Authentication and EAP-TLS Authentication	Supports both Machine Certificate Authentication and EAP-TLS Authentication	Supports both Machine Certificate Authentication and EAP-TLS Authentication	Supports both Machine Certificate Authentication and EAP-TLS Authentication (Profile Configuration can be customized to use certificate or EAP - TLS)

Comparison Parameter	Windows Desktop/Laptop			Windows Mobile Phone		Linux Client	iOS Client	MAC OS Client
	Machin e Certi ficate Auth entica tion Certi ficate EKU Exten sion	Client Certifica te should have cli- en- tAuth (1.3.6.1. 5.5.7.3. 2) and serv- erAuth (1.3.6.1. 5.5.7.3. 1) EKU Extensio n mandat o-rily. Op- tionally Certi- ficates can have Se-cure Email (1.3.6.1. 5.5.7.3. 4) or Encrypt- ing File System (1.3.6.1. 4.1.311. 10.3.4)	Client Cer- tificate should have clien- tAuth (1.3.6.1. 5.5.7.3. 2) and serv- erAuth (1.3.6.1. 5.5.7.3. 1) EKU Ex- tension mandat o-ri-ly. Optiona l-ly Certi- ficates can have Secure Email (1.3.6.1. 5.5.7.3. 4) or Encrypti ng File System (1.3.6.1. 4.1.311. 10.3.4)	Client Certifica te should have cli- en- tAuth (1.3.6.1. 5.5.7.3. 2) and serv- erAuth (1.3.6.1. 5.5.7.3. 1) EKU Ex- tension manda- torily. Optiona lly Certi- ficates can have Secure Email (1.3.6.1. 5.5.7.3. 4) or Encrypti ng File System (1.3.6.1. 4.1.311. 10.3.4)	Not Appli- cable	Not Appli- cable	Client Cer- tificate should have cli- en- tAuth (1.3.6.1. 5.5.7.3. 2) and serv- erAuth (1.3.6.1. 5.5.7.3. 1) EKU Extensio n mandat o-rily. Op- tionally Certi- ficates can have Se-cure Email (1.3.6.1. 5.5.7.3. 4) or Encrypti ng File Sys-tem (1.3.6.1. 4.1.311. 10.3.4)	Client Certifi- cate should have clien- tAuth (1.3.6.1. 5.5.7.3. 2) and serv- erAuth (1.3.6.1. 5.5.7.3. 1) EKU Extensio n mandat o-ri-ly. Optiona lly Cer- tificates can have Secure Email (1.3.6.1. 5.5.7.3. 4) or Encrypt- ing File System (1.3.6.1. 4.1.311. 10.3.4)

Comparison Parameter	Windows Desktop/Laptop			Windows Mobile Phone		Linux Client	iOS Client	MAC OS Client
DH 2048 bit or DH 3072 bit for Phase 1 key negotiation	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported	Supported	Supported	Supported
DH Support for Phase 1 key negotiation	DH1024	DH1024	DH1024	DH1024	DH1024	DH1024 DH2048 DH3072	DH1024 DH2048 DH3072	DH2048

Comparison Parameter	Windows Desktop/Laptop			Windows Mobile Phone		Linux Client	iOS Client	MAC OS Client
	ICS	ICS	SAs	ICS	ICS			
SA (Security Association) Preference Order	ICS chooses matching SAs based on the Preference Order sent by IKEv2 Client. Note: ICS doesn't maintain any default SA Preference Order	ICS chooses matching SAs based on the Preference Order sent by IKEv2 Client. Note: ICS doesn't maintain any default SA Preference Order	SAs based on the Preference Order sent by IKEv2 Client. ICS doesn't maintain any default SA Preference Order	ICS chooses matching SAs based on the Preference Order sent by IKEv2 Client. Note: ICS doesn't maintain any default SA Preference Order	ICS chooses matching SAs based on the Preference Order sent by IKEv2 Client. ICS doesn't maintain any default SA Preference Order	ICS chooses matching SAs based on the Preference Order sent by IKEv2 Client. ICS doesn't maintain any default SA Preference Order	ICS chooses matching SAs based on the Preference Order sent by IKEv2 Client. ICS doesn't maintain any default SA Preference Order	ICS chooses matching SAs based on the Preference Order sent by IKEv2 Client. ICS doesn't maintain any default SA Preference Order

- Windows IKEv2 Native Client doesn't Support DH2048 and above, so on Enabling 'Allow only DH 2048 bit and higher for Phase 1 key negotiation' Checkbox IKEv2 Negotiation will fail.
- DH1536, DH768, DH4096 and Higher Diffie Hellman Algorithms are not Supported. Currently ICS Supports only following Diffie Hellman Algorithms
 - DH1024
 - DH2048
 - DH3072

- ICS doesn't enforce SA (Security Association) Preference Order in IKEv2 Phase1 Negotiation, ICS only honors the SA Preference Order what IKEv2 Client Sends.
- IKEv2 Configuration doesn't Support Port/Realm Mapping for the Virtual Ports having same name Un-der Internal and External Ports.
- IKEv2 Client doesn't Support Host Checker Validation both at Realm and Role Level.
- IKEv2 in ICS doesn't support IPv6 Address.
- IKEv2 Client doesn't honor Roaming Session Settings under Roles Session Options
- Due to Design Limitation following system operation is not supported for IKEv2 Configuration
 - XML Export from a ICS running 8.2Rx build and Import to another ICS running 8.3R1
 - Binary Export from a ICS running 8.2Rx build and Import to another ICS running 8.3R1
 - Push Config (Selective Config) from a ICS running 8.2Rx build to another ICS running 8.3R1
 - Push Config (Entire Config) from a ICS running 8.2Rx build to another ICS running 8.3R1
 - Pulse One doesn't Support Pushing IKEv2 Configuration
- IKEv2 does not support automatic cluster failover. After cluster failover, IKEv2 users must reconnect.
- IKEv2 clients do not Support IPSEC negotiation with ECC device certificate configured in ICS.
- AES256/MD5 and AES128/MD5 ESP Encryption is not Supported by Windows Native Client and Mobile Phone.
- VPN Tunneling Connection Profile Proxy Server Settings under Users -> Resource Policies -> VPN Tun-eling Connection Profiles is not Supported by IKEv2 Clients.
- Windows 10 VPN Client Proxy does not work with ICS.
- Windows 10 Native Client or Windows 10 Mobile does not use or support split tunnel configuration of ICS for routing Traffic.
- Deny/Exclude Access in Split Tunnel Network Profile Configuration doesn't work with IKEv2 Clients
- IKEv2 Native Clients won't honor "Key lifetime (time based)" and "Key lifetime (bytes transferred)" Connection Profile Configuration in ICS for IPSEC SA Rekeying
- MAC OS 10.12 IKEv2 Client will automatically disconnects after 8 minutes

Configuring IKEv2 Ports

To configure the IKEv2 ports and EAP protocol:

1. Select **System > Configuration > IKEv2** to display the configuration page. See Figure underneath.
2. Enter the DPD timeout value in seconds. Valid values are 400-3600.

DPD is a form of keepalive. When a tunnel is established but idle, one or both sides may send a "hello" message and the other replies with an acknowledgement. If no response is received, this continues until the DPD time value has elapsed. If there still isn't any traffic or acknowledgement, the peer is de-termined to be dead and the tunnel is closed.

3. Under Port/Realm Mapping, select the port and the realm to use that port.
 - To add additional port/realm mapping sets, click **Add**.
 - To delete a port/realm mapping set, select the check box next to the set to remove and click Delete.
4. Under Realm / Protocol Set Mapping, select the realm and the EAP protocol set to use for that realm. The three Protocol Set Options include **EAP-MSCHAP-V2, EAP-MD5-Challenge, and EAP-TLS**.
 - To add additional realm/protocol mapping sets, click **Add**.
 - To delete a realm/protocol mapping set, select the check box next to the set to remove and click **De-lete**.
5. Click **Save Changes**.



Changing IKEv2 configuration (System > Configuration > IKEv2) disconnects connections from IKEv2 clients, VPN Tunneling and Ivanti. VPN Tunneling and Ivanti will reconnect automatically.

The following figure depicts the IKEv2 Configuration for EAP-TLS:

DPD Timeout: Dead peer detection timeout (400-3600 seconds)

Port / Realm Mapping

Specify the ports IKEv2 users can connect to and their associated user authentication realm

[Delete](#)

<input type="checkbox"/>	Port	Realm	
<input type="checkbox"/>	<input type="text" value="internal (10.96.144.49)"/> ▼	<input type="text" value="Users"/> ▼	Add
<input type="checkbox"/>	internal (10.96.144.49)	Users	

Realm / Protocol Set Mapping

Specify the EAP protocol sets that are to be supported by the user realms.
Note that user realms using certificate authentication for IKEv2 users do not need to be associated with any protocol sets.

[Delete](#)

<input type="checkbox"/>	Realm	Protocol set	
<input type="checkbox"/>	<input type="text" value="Users"/> ▼	<input type="text" value="EAP-MSCHAP-V2"/> ▼	Add
<input type="checkbox"/>	Users	EAP-MSCHAP-V2	

Phase 1 key settings

IKEv2 Phase 1 key settings below will override current key settings. Changing these settings will disconnect and reconnect clients

Allow only AES256 for Phase 1 key negotiation

Allow only SHA2 for Phase 1 key negotiation

Allow only DH 2048 bit or DH 3072 bit for Phase 1 key negotiation

Initial Contact

Enabling Initial Contact will delete all existing sessions for that user if request contains INITIAL_CONTACT payload when Multi user session is enabled

Enable PDS to process INITIAL_CONTACT request

[Save changes?](#)

To enable IKEv2 User Access Logs:

1. Select **System > Logging/Monitoring > User Access > Log Settings**.
2. Under Select Events to Log, make sure to enable the **Pulse Client Messages** checkbox.
3. Click **Save Changes**.

IKEv2 Configuration Overview

IKEv2 EAP supports the following authentication server types:

- Local authentication

- Active Directory
- Certificate Server (applicable only for EAP-TLS)

If you are using IKEv2 EAP authentication on a local authentication server, you must select the Pass-word stored as clear text check box in the Auth Server page of the admin console. Note that you can-not edit an existing local authentication server instance to select this option. If you require IKEv2 EAP authentication on a local authentication server, you must create a new local authentication server in-stance.



IKEv2 EAP does not work with any preexisting local authentication servers since they do not store passwords in clear text.

To configure support for IKEv2:

1. Configure your client for using IKE. For more information, see your mobile device's documentation.
2. Install client and device certificates.
 1. You need a Certificate Authority (CA) that can issue client certificates.
 2. On the client side, install this client certificate along with the CA certificate.
 3. On the Ivanti Connect Secure server side, install the CA certificate under Configuration/Certificates/Trusted Client CAs.
 4. On the client side, install the Connect Secure certificate corresponding to the port to which the client connects, found under Configuration/Certificates/Device Certificates.
3. Define an IKEv2 rule under the Users > User Realms > User > Role Mapping page of the admin console.
4. Select the IKEv2 access feature under the Users > User Roles > User > General > Overview page of the admin console.
5. Enable Network Connect for the Role and configure an NC Connection Profile (IP pool) to use for that Role.

When a client uses IKEv2 to connect to the host, the Agent Type column of the Active Users page displays IKEv2.

Enabling the IKEv2 Phase-1 Key Settings

IKEv2 has a two-phase negotiation process. The first phase is known as IKE_SA_INIT and the second phase is known as IKE_AUTH. IKE_SA_INIT Phase exchanges the Security Association (SA) proposals, which comprises Encryption and Integrity algorithms, Diffie-Hellman Group, to derive Keys for IKE_AUTH Phase. These Security Association proposals preference can be controlled by different configurations in ICS.

To Configure Phase-1 Key Settings, select System > Configuration > IKEv2 > Phase 1 Key Settings. Three new UI options are available to enforce Encryption Algorithm (AES256), Integrity Algorithm (SHA256, SHA384 and SHA512) and Diffie-Hellman Group (DH 2048 and DH3072). Enabling these options mean more secured Phase 2 negotiations. When AES256 is enabled, AES256 Encryption Algorithm is preferred over AES128 or 3DES. When SHA2 is Enabled, SHA2 Integrity Algorithm is preferred over SHA1 and When DH is Enabled, DH2048 or DH3072 Diffie-Hellman Group is preferred over DH1024. See figure below for Phase-1 Key Settings (In the main Menu bar: **System**):

Specify the EAP protocol sets that are to be supported by the user realms.
Note that user realms using certificate authentication for IKEv2 users do not need to be associated with any protocol sets.

Delete

Realm	Protocol set	
Users	EAP-MSCHAP-V2	Add
AD-Realm	EAP-MSCHAP-V2	
Cert-Realm	EAP-TLS	

Phase 1 key settings
IKEv2 Phase 1 key settings below will override current key settings Changing these settings will disconnect and reconnect clients.

Allow only AES256 for Phase 1 key negotiation

Allow only SHA2 for Phase 1 key negotiation

Allow only DH 2048 bit or DH 3072 bit for Phase 1 key negotiation

Save changes?

Save Changes

By default, these check boxes are disabled for backward compatibility. Enabling the check boxes will override current key settings and will disconnect connected clients if any.

Configuring IKEv2 Phase-2 Key Settings

The Phase-2 Key Exchange is also known as IKE_AUTH. The IKE_AUTH exchange is used to authenticate the remote endpoint and to securely establish IPSec Security association or Child SA. Only PSA platforms support SHA256. The following encryption and integrity combinations are supported:

- AES128 + SHA1/MD5
- AES256 + SHA1/MD5/SHA256
- AES128 + SHA1/MD5

To configure IKEv2 Phase-2 parameters:

1. Select **Resource Policies -> VPN Tunneling -> Connection Profile**.
2. Under **Encryption**, select the suitable encryption and integrity combination.

The image below, indicates the options available for encryption:



The image shows a configuration window for encryption. On the left, the word "Encryption:" is followed by a large empty text box. On the right, there is a list of five radio button options:

- AES128/MD5 (MD5 is insecure. Option is not recommended)
- AES128/SHA1
- AES256/MD5 (MD5 is insecure. Option is not recommended)
- AES256/SHA1
- AES256/SHA256 (maximize security)

Enabling the IKEv2 Initial Contact

When an endpoint either crashes or reinitializes its state, the other endpoint should detect those conditions and stop sending any data. The INITIAL_CONTACT notification asserts that IKE Security Association (SA) is the only IKE SA currently active between the authenticated identities. It may be sent when an IKE SA is established after a crash, and the recipient may use this information to delete any other IKE SAs it has to the same authenticated identity without waiting for a timeout.

Enabling Initial Contact deletes all existing sessions for that user if request contains INITIAL_CONTACT payload when Multi user session is enabled.



When multiuser session is disabled, the server will always delete the existing session for that user before creating a new session.

To configure IKEv2 Initial Contact:

1. Select **System > Configuration > IKEv2**.
2. In the Initial Contact section, select the **Enable ICS to process INITIAL_CONTACT** request check box.

Enabling the IKEv2 Access Feature

Roles specify the session properties, including enabled access features, for users who are mapped to the role.

To enable the IKEv2 access feature:

1. Select **User > User Roles > Role Name > General > Overview**.
2. Under **Access Features**, check the **VPN Tunneling** check box.
3. Click **Save Changes**.

Enabling the IKEv2 EAP TLS User Access Logs

1. Select **System > Logging/Monitoring > User Access > Log Settings**.
2. Under **Select Events** to Log, make sure to enable the **Pulse Client Messages** check box.
3. Click **Save Changes**.

Defining the IKEv2 Role Mapping Rule

Role mapping rules are conditions a user must meet for the system to map the user to one or more user roles.



The procedure described in this topic is required only if you want to create a separate role mapping rule specific for IKEv2 users. If you use regular username, group or custom expression-based role mapping rules (typically used for general access to a device), the following procedure is not required.

1. Select **User > User Realms > User > Role Mapping**.
2. Click **New Rule**.
3. Select Custom Expressions as the type of condition on which to base the rule.
4. Click **Update** to display the Expressions list.
5. Click the Expressions button to display the Expressions tab of the server catalog.
6. Create a rule: **userAgent = "IKEv2"**.
7. Click **Add Expression** and then **Close**.
8. Select the rule you just created from the Available Expressions list and click **Add** to move it to the Selected Expressions list.
9. Specify the roles to assign to the authenticated user by adding roles to the Selected Roles list.
10. (optional) Check the Stop processing rules when this rule matches check box if you want the system to stop evaluating role mapping rules when the user meets the conditions specified for this role.
11. Click **Save Changes**.

Using the Mobile Options

This topic describes the mobile options that are available on Ivanti Connect Secure. To configure the mobile option, go to System > Configuration > Mobile. It includes the following information:

The screenshot shows the 'Configuration' page for 'Mobile' settings. The navigation bar includes tabs for Licensing, Pulse One, Security, Certificates, DMI Agent, NCP, Client Types, Virtual Desktops, User Record Synchronization, IKEv2, SAML, and Mobile (selected). Below the navigation bar, there are three main configuration sections:

- Server certificate trust enforcement**: Enabling this option would block connections if the Ivanti Connect Secure server certificate is untrusted or invalid.
 - Enabled: Automatically block Pulse Secure mobile app from connecting to untrusted Ivanti Connect Secure.
 - Disabled: Prompt the user when Pulse Secure mobile app connects to untrusted Ivanti Connect Secure.
- Reconnect VPN on wakeup**: Enabling this option would reconnect VPN session with PCS on device wakeup.
 - Enabled: Reconnect VPN session on wakeup.
 - Disabled: Do not reconnect VPN session on wakeup.
- Touch ID / Face ID Support for iOS devices**: Enabling this option would allow users to authenticate using fingerprint / face recognition.
 - Enabled: Enable Touch ID / Face ID for user authentication.
 - Disabled: Disable Touch ID / Face ID for user authentication.

A 'Save Changes' button is located at the bottom left of the configuration area.

The following table lists the Configuring the Mobile Options:

Option	Description
Server certificate trust enforcement	Enables you to block connections if the Ivanti Connect Secure server certificate is untrusted or invalid. When enabled, it automatically blocks the Pulse mobile app from connecting to untrusted Ivanti Connect Secure. When disabled, it prompts when a Pulse mobile app connects to untrusted Ivanti Connect Secure.
Reconnect VPN on wakeup	Enables you to reconnect a VPN session with ICS on device wakeup.
Touch ID / Face ID Support for iOS devices	Enables you to authenticate using fingerprint / face recognition.

Using the Advanced Client Configuration Feature

This topic describes the XML advanced client configuration that can be used by the ICS administrator to configure the custom settings, which are meant to solve a specific customer scenario without changing the ICS admin console. Admin can set these custom settings in the form of XML input through the Advanced Client Configuration UI feature. Pulse clients supporting these custom settings will consume them when connecting to this ICS, and the same would be applied on the client machines. From 9.0R3 release onwards, this feature will minimize the number of changes going into the ICS admin console, in order to fulfill a custom requirement of a specific customer.

In the earlier Pulse client releases, i.e. prior to v5.2R2, the virtual adapter MTU was calculated based on the physical adapter MTU (of the host machine) and the MTU sent by the ICS.

Basically, the formula used to calculate the virtual adapter MTU is:

MIN (Physical Adapter MTU, MTU from ICS, TCP MSS value + 40)

Following is one of the scenarios where Firewall on the data path is stripping the TCP MSS options being advertised by ICS to the Pulse client. In this scenario, the TCP MSS value on the Pulse client will default to a minimum value of 536, and as a result the client side MTU calculation will result in a minimum MTU value of 576. Here, customer wants to ignore the TCP MSS options while calculating the Virtual Adapter MTU calculation.

If the administrator configures the Ivanti Connect Secure sever with the following XML input in "Advanced Client Configuration for Pulse Client" option, it will ignore TCP MSS options while calculating the virtual adapter MTU on client side.

1. Select **System > Configuration > Advanced Client Configuration** to display the configuration page. Figure shows the configuration page for Ivanti Connect Secure.

The screenshot shows the configuration page for Ivanti Connect Secure. The breadcrumb navigation is "Configuration > Advanced Client Configuration". The page title is "Advanced Client Configuration". There is a horizontal menu with various options: Licensing, Pulse One, Security, Certificates, DMI Agent, NCP, Client Types, Virtual Desktops, User Record Synchronization, IKEv2, SAML, Mobile, VPN Tunneling, PSAM, Telemetry, **Advanced Client Configuration** (highlighted), and Advanced Networking. Below the menu, there is a section titled "Advanced Configuration for Pulse clients" with a sub-header "Enter advanced configuration in XML". A large text area is provided for entering XML configuration. At the bottom left of the text area is a "Save Changes" button.

2. Enter the following XML input in "Advanced Configuration for Pulse Clients".

```
<advanced-config>
  <version>9.0.3</version>
  <desktop-client-config>
    <layer3-connection-config>
      <adapter-config>
        <ignore-tcp-mss>TRUE</ignore-tcp-mss>
      </adapter-config>
    </layer3-connection-config>
  </desktop-client-config>
</advanced-config>
```

3. Click **Save Changes**.

The advanced configuration setting "ignore-tcp-mss" is Layer3 Adapter configuration setting and this will be consumed by the Pulse client as part of the IpsecConfig.



This "ignore-tcp-mss" setting is applicable for the virtual adapter MTU calculation only for IPv4. By default, the setting is always false, and therefore the TCP MSS options are always considered for MTU by default. Admin has to explicitly set the ignore-tcp-mss setting to TRUE (case-insensitive), to ignore TCP MSS.

Using the Traffic Segregation Feature

This topic describes the traffic segregation feature that is available on the Connect Secure virtual appliance (service provider edition).

Traffic Segregation Feature Overview

Service providers often need a way to cleanly segregate the system-generated network traffic across interfaces (such as Internal, Management and VLAN ports), to differentiate AAA traffic from others.

Traffic segregation is supported for the following Scenarios:Radius

Certificate Auth including ANY CRL/OCSP verification.

- SAML
- AAA DNS Traffic
- DMI
- System logging (syslog)
- AD- Domain Join

- AD- Server Catalog
- AD-User Auth
- AD-Authrz
- AD-PMI
- LDAP-Test Connection
- LDAP-User Auth
- LDAP-User Auth- Referral user
- LDAP-SearchCatalog
- LDAP-Grplookup-UserLogin
- LDAP-PMI

Unsupported features include the following:

- Ace Auth



- For 9.0R2 and previous releases, enable the Send AAA Traffic via Management Port to send AAA traffic through management port. From 9.0R3 release, this option is enhanced and modified. For more information see, [AAA Traffic Management](#).
- DMI/Netconf listens only on the default internal IP and default management IP. VIP IP addresses must not be used for DMI communication with ICS.

Two typical service provider deployment models are:

- Authentication server of the customer is hosted by the customer
- Authentication server for the customer is hosted and managed by the service provider

In both models, the service provider's authentication server is always hosted in the service provider's network and is reachable either through the internal or management port. In the first model, the customer's authentication servers are reachable through the internal port of the virtual appliance. In the second model, the customer's authentication server must be routed either through the internal or management port, depending on where the service provider has hosted the customer's authentication server.

A Traffic Segregation menu is available only on virtual appliances to allow system providers to configure the interface and traffic. The "Default Network" is used as the primary logical network for the service provider and customer configuration. If traffic segregation across different logical networks is needed, the "Administrative Network" can be used.

You can differentiate AAA traffic from other traffic and route it through the management port. This is useful when the only the authentication servers are hosted on the network reachable through the management port and all other resources uses a different port. This option is available on both the Default Network and the Administrative Network.

The configurations to do on the virtual appliance depend on the logical network setup around the virtual appliance and the type of service provider deployment model:

- If both the service provider's and customer's authentication server are reachable through the same interface, the entire configuration for the service provider and customer is done under the Default Network. It is not necessary to enable the Administrative Network.
- If the service provider's and customer's authentication servers are located on two different networks, the Administrative Network must be created. The following table shows where the administrator configures the options in the system GUI.

The following table describes the Configuring Traffic Segregation Options:

Options	Description
Network Settings and Tools	Enables you to change standard network settings; print a routing table; print or clear an ARP cache; run the ping and traceroute commands, remove static routes, add an ARP entry, view cluster status, configure management port, and change traffic control settings (Note: For change traffic control settings, the goal of the change is to change the priority of traffic in IVE depending on its criticality).
Create admin username and password	Enables you to create a new super administrator account.
Display log	Enables you to display system configuration, user access logs, or administrator access logs through the serial console. Note that must enter q to return to serial console options after viewing the logs.
System Operations	Enables you to reboot, shut down, restart, roll back, or factory reset the system without using the admin console.

Options	Description
Toggle password protection for the console	Enables you to password protect the serial console. When you toggle this option to "on," only super administrators are allowed access.
Create a Super Admin session	<p>Enables you to create a recovery session to the admin console, even if you have configured the system to block access to all administrators. When you select this option, the system generates a temporary token that is valid for 3 minutes. Enter the following URL into a browser window:</p> <p>Then, enter the temporary token when prompted to sign in to the admin console.</p> <p>When you select this option, the system blocks any additional administrators from signing in to the admin console until you sign in to the specified URL and initiate a session using your token. The appliance blocks additional sign-in attempts so that you can fix any configuration problems that the system may have encountered without conflicting with another session.</p>
System Snapshot	<p>Enables you to take a system snapshot without using the admin console. When you select this option, the system takes the snapshot immediately. You can then send the snapshot file, by way of SCP, to a remote system. The system prompts you for the destination server port, user ID, password, and the destination path to the remote directory.</p> <p>If you choose not to send the snapshot file to a remote system, the system saves the file locally. The next time you log in to the admin console, the System Snapshot tab contains a link to the snapshot file.</p>

Using the Serial Port

This topic describes use of the serial port and serial port console.

Connecting to the Serial Port Console

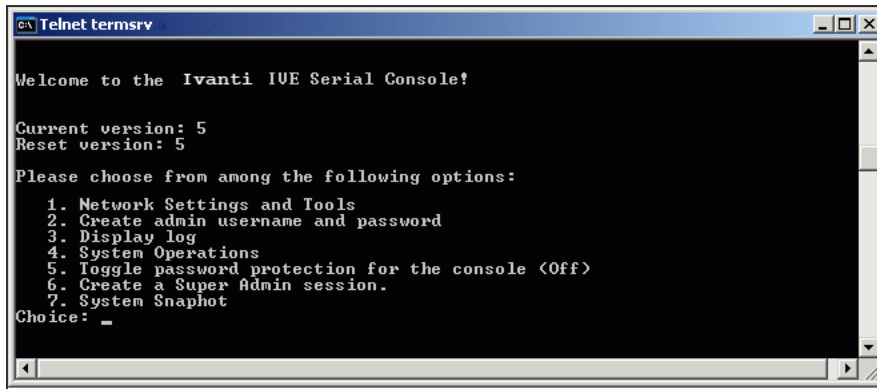
In cases where the admin console is unavailable, you can perform network and host configuration tasks and troubleshooting using the serial port console.

To connect to the serial console:

1. Plug a null modem crossover cable from a console terminal or laptop into the device serial port. This cable is provided in the product box. Do not use a straight serial cable.

2. Configure a terminal emulation utility, such as HyperTerminal, with the following serial connection pa-rameters:
 - 9600 bits per second
 - 8-bit No Parity (8N1)
 - 1 Stop Bit
 - No flow control
3. Press **Enter** until the serial console is displayed.

The following figure depicts the Serial Console Menu Options:



The following table lists the Serial Console Menu:

Options	Description
Network Set-tings and Tools	Enables you to change standard network settings; print a routing table; print or clear an ARP cache; run the ping and traceroute commands, remove static routes, add an ARP entry, view cluster status, configure management port, and change traffic control settings (Note: For change traffic control settings, the goal of the change is to change the priority of traffic in IVE depending on its criticality).
Create admin username and password	Enables you to create a new super administrator account.
Display log	Enables you to display system configuration, user access logs, or administra-tor access logs through the serial console. Note that must enter q to return to serial console options after viewing the logs.

Options	Description
System Operations	Enables you to reboot, shut down, restart, roll back, or factory reset the system without using the admin console.
Toggle password protection for the console	Enables you to password protect the serial console. When you toggle this option to "on," only super administrators are allowed access.
Create a Super Admin session	Enables you to create a recovery session to the admin console, even if you have configured the system to block access to all administrators. When you select this option, the system generates a temporary token that is valid for 3 minutes. Enter the following URL into a browser window: Then, enter the temporary token when prompted to sign in to the admin console. When you select this option, the system blocks any additional administrators from signing in to the admin console until you sign in to the specified URL and initiate a session using your token. The appliance blocks additional sign-in attempts so that you can fix any configuration problems that the system may have encountered without conflicting with another session.
System Snapshot	Enables you to take a system snapshot without using the admin console. When you select this option, the system takes the snapshot immediately. You can then send the snapshot file, by way of SCP, to a remote system. The system prompts you for the destination server port, user ID, password, and the destination path to the remote directory. If you choose not to send the snapshot file to a remote system, the system saves the file locally. The next time you log in to the admin console, the System Snapshot tab contains a link to the snapshot file.

Using the Serial Console to Roll Back to a Previous OS Version

You can use the admin console to roll back the configuration to a previous state. If the rollback option is not available in the admin console, you can use the procedure described in this section to perform the system rollback.

If you have not yet performed an OS service package upgrade, there is no previous state to roll back to, and the rollback option is not available. If you have performed an OS service package upgrade, any system and user configuration data created after the upgrade is lost unless you export the most current configuration files before rolling back the system and then import them afterwards.

To roll back to the previous OS service package:

1. Connect to the serial console.
2. In a browser window, sign in to the admin console.
3. Select **Maintenance > System > Platform**.
4. Click **Reboot Now** and then return to the console utility window. The window displays a message that the system is restarting.
5. After several moments, you are prompted to use the Tab key to select options. Press Tab, and when prompted for the configuration to load, type rollback and then press Enter.

After you click **Reboot Now**, the rollback status is output to the screen, and when complete, you are prompted to press Return (Enter) to modify system settings, which returns you to the initial setup options. When you are finished entering data, simply close the serial console window.

If you wait more than 5 seconds to enter your choice, the current system configuration is automatically loaded, and you must go back to the admin console and click Reboot Now to start the process again. If you have already performed a system rollback, the rollback option is not available again until you upgrade the OS service package again.

Using the Serial Console to Reset the System to the Factory Image

In rare cases, you might need to reset the system to its original factory settings. Before performing this advanced system recovery option, contact Global Support Center (<https://www.pulsesecure.net/support/>). If possible, export the most current system and user configuration data before performing a factory reset.

To perform a factory reset:

1. Connect to the serial console.
2. In a browser window, sign in to the admin console.
3. Select **Maintenance > System > Platform**.
4. Click **Reboot** and then go back to the console utility window. The window displays a message that the system is restarting.
5. After several moments, you are prompted to use the Tab key to select options. Press Tab, and when prompted for the configuration to load, type factory-reset and then press Enter.

If you wait more than 5 seconds to enter your choice, the current system configuration is automatically loaded, and you must go back to the admin console and click **Reboot Now** to start the process again.

6. When you are prompted to confirm performing a factory reset, type `proceed` and then press `Enter`.

The system begins the process of resetting the machine to its original settings and outputs several screens of data. After several minutes, you are prompted to use the `Tab` key to select configuration choices.

7. When prompted to press the `Tab` key, do one of the following:

1. Wait for the default selection (`current`) to start automatically.
2. Press **Tab**, type `current`, and then press **Enter**.

You are then prompted to enter the initial configuration settings. For details on how to proceed, see the Installation Guide provided in the product packaging or on the Ivanti Support site.

After you complete the initialization process, you can upgrade to the latest OS service package and import saved system and user configuration files to return to the last good working state of your system.

You might receive errors from the system during the initial setup or on a factory reset. Before the system starts services, it monitors the network port for a maximum of 120 seconds. The system checks the link status and sends ARP requests to the default gateway. If there is a problem, after 5 seconds, the system displays a message on the serial console that starts with **NIC:.....** If the link recovers within 120 seconds, the startup process continues. If the link does not recover, the following message is displayed:

```
Internal NIC: ..... [Down code=0x1]
```

Two codes can appear:

- **0x1** means that the interface link status reported by the NIC remains off (for example, a disconnected cable or a cable is in the wrong port).
- **0x2** means that the gateway is unreachable. The system boots but is not reachable from IP addresses bound to that network port.

Certificate Security Administration

Understanding Digital Certificate Security

Ivanti Connect Secure uses Public Key Infrastructure (PKI) to secure the data sent to clients over the Internet. PKI is a security method that uses public and private keys to encrypt and decrypt information. These keys are enabled and stored through digital certificates. A digital certificate is an encrypted electronic file issued by a certificate authority (CA) that establishes credentials for client/server transactions.

In public key cryptography, a public/private key pair is used to encrypt and decrypt data. Data encrypted with a public key, which the owner makes available to the public, can be decrypted with the corresponding private key only, which the owner keeps secret and protected. For example, if User1 wants to send User2 an encrypted message, User1 can encrypt it with User2's public key and send it. User2 then decrypts the message with the private key. The reverse process is also useful: encrypting data with a private key and decrypting it with the corresponding public key. This process is known as creating a digital signature. For example, if User1 wants to present User1's own identity as the sender of a message, User1 can encrypt the message with User1's private key and send the message to User2. User2 then decrypts the message with User1's public key, thus verifying that User1 is indeed the sender.

Ivanti Connect Secure systems use the following types of digital certificates to establish credentials and secure session transactions:

- Device certificates-A device certificate helps to secure network traffic to and from the Pulse client service using elements such as company name, a copy of your company's public key, the digital signature of the CA that issued the certificate, a serial number, and expiration date.
- Trusted client CAs-A trusted client CA is a CA that issues client-side certificates. You can use trusted client CAs in the access management framework realm and role configurations to require certificates or certificates with specific attributes. For example, you may specify that users must present a valid client-side certificate with the OU attribute set to "yourcompany.com" to sign into the Users authentication realm.
- Trusted server CAs-A trusted server CA is a CA which issues certificates for web server. You can install a trusted server CA to validate the credentials of the web sites that users access through the Pulse client service.
- Code-signing certificates-A code-signing certificate (also called an applet certificate) is a certificate that re-signs Java applets that are intermediated by Ivanti Connect Secure. You can use the self-signed code-signing certificate that comes pre-loaded, or you can install your own code-signing certificate.

- Client auth certificates-In this context, the client auth certificate is used when backend SSL servers require Ivanti Connect Secure to present a client certificate for authentication.



- The system can verify certificates that use SHA2 as the message digest.
 - DSA certificates are not supported.
-

Using Device Certificates

This topic describes how to use device certificates.

Understanding Device Certificates

A device certificate helps to secure network traffic to and from the Pulse client service using elements such as your company name, a copy of your company's public key, the digital signature of the Certificate Authority (CA) that issued the certificate, a serial number, and an expiration date.

When receiving the device certificate from the system, the client's browser first verifies whether the device certificate is valid and whether the user trusts the CA that issued the certificate. If the user has not already indicated that they trust the certificate issuer, the Web browser prompts the user with a warning saying connection is untrusted or there is a problem with the websites security certificate.


The system supports X.509 device certificates in DER and PEM encode formats (file extensions include .cer, .crt, .der, and .pem) as well as PKCS #12 (file extensions include .pfx and .p12). The system also supports the following features:

- Intermediate device CA certificates-Within a certificate hierarchy, one or more intermediate certificates are branched off a single root certificate.
- Multiple device certificates-When using multiple device certificates, each certificate handles validation for a separate hostname or fully qualified domain name (FQDN) and can be issued by a different CA.

Understanding Self-Signed Certificates

When you initialize the system with the serial console, the system creates a self-signed certificate that enables you to immediately begin setting up the system. Users are prompted with a security alert each time they sign in because the certificate is not issued by a trusted CA. Figure 193 shows the security alert.

The following figure depicts the Security Alert When the Device Certificate Is Not Issued by a Trusted CA:



There is a problem with this website's security certificate.

The security certificate presented by this website was not issued by a trusted certificate authority. The security certificate presented by this website was issued for a different website's address.

Security certificate problems may indicate an attempt to fool you or intercept any data you send to the server.

We recommend that you close this webpage and do not continue to this website.

- [Click here to close this webpage.](#)
- [Continue to this website \(not recommended\).](#)
- [More information](#)

If you use the self-signed certificate created upon initialization with the serial console, users are prompted with a certificate warning when accessing the sign-in page.

Before promoting the system to production use, we recommend you replace the self-signed certificate with a certificate issued by a trusted CA.

Importing a Device Certificate and Private Key

The system uses certificates to verify itself to other network devices. A digital certificate is an electronic means of verifying your identity through a trusted third party, known as a Certificate Authority (CA). Your company might use its own enterprise CA server, or it might use a reputable third-party CA.

To import an enterprise root server certificate and private key:

1. Select System > Configuration > Certificates > Device Certificates.
2. Click Import Certificate & Key to display the configuration page.
3. One of the following options to complete the import procedure:
 - If certificate file includes private key-When the certificate and key are contained in one file.
 - If certificate and private key are separate files-When the certificate and key are in separate files.
 - Import via System Configuration file-When the certificate and key are contained in a system configuration file. With this option, the system imports all of the certificates specified (including private keys and pending CSRs, but not the corresponding port mappings).

In the appropriate form, browse to the certificate and key files. If the file is password-protected, enter the password key.

4. Click Import.

Creating a Certificate Signing Request

If your company does not own a digital certificate for its Web servers, you can create a certificate signing request (CSR) and then send the request to a CA for processing. When you create a CSR, a private key is created locally that corresponds to the CSR. If you delete the CSR at any point, this file is also deleted, prohibiting you from installing a signed certificate generated from the CSR.

To create a certificate signing request:

1. Select System > Configuration > Certificates > Device Certificates.
2. Click New CSR to display the configuration page.
3. Complete the required information and click Create CSR.
4. Follow the onscreen instructions, which explain what information to send to the CA and how to send it.

When you submit a CSR to a CA authority, you might be asked to specify either the type of Web server on which the certificate was created or the type of Web server the certificate is for. Select apache (if more than one option with apache is available, select any). If you are prompted for the certificate format to download, select the standard format.

Do not send more than one CSR to a CA at one time. Doing so can result in duplicate charges.



- To view details of any pending requests that you previously submitted, click the Certificate Signing Request Details link.
 - While generating a CSR, an apostrophe string is required, prefix it by an escape character. For example, "Children's" should be "Children\'s".
-

Importing a Signed Certificate Created from a CSR

When you receive the signed certificate from the CA, import it.

To import a signed device certificate created from a CSR:

1. Select System > Configuration > Certificates > Device Certificates.

2. Under Certificate Signing Requests, click the Pending CSR link that corresponds to the signed certificate.
3. Under Import signed certificate, browse and select the certificate file you received from the CA, and then click Import.

Understanding Intermediate Certificates

Within a certificate hierarchy, one or more intermediate certificates are branched off a single root certificate. The root certificate is issued by a root CA and is self-signed. Each intermediate certificate is issued by the certificate preceding it in the chain.

To use chained certificates in your deployment, you must ensure that the server (Ivanti Connect Secure) and client (Web browser) together contain the entire certificate chain. For example, you can secure traffic using a chain that stems from a VeriSign root certificate. If your users' browsers come preloaded with VeriSign root certificates, you need to install only the lower-level certificates in the chain. When your users sign in, the system presents any required certificates within the chain to the browser to secure the transaction. The system creates the proper links in the chain using the root certificate's IssuerDN. If the system and browser together do not contain the entire chain, the user's browser does not recognize or trust the device certificate because it is issued by another certificate instead of by a trusted CA.

You can upload one or more intermediate CAs in a PEM file. The entire chain must be sent to the client in descending order, starting with the root certificate.

Within a certificate hierarchy, one or more intermediate certificates are branched off a single root certificate. The root certificate is issued by a root CA and is self-signed. Each intermediate certificate is issued by the certificate preceding it in the chain.

To use chained certificates in your deployment, you must install the appropriate client-side certificates in each user's Web browser and then upload the corresponding CA certificates to Ivanti gateway's Intermediate CA store. Use one of the following methods to upload the certificate chain:

- Import the entire certificate chain in one file. The file must contain the root certificate and any subcertificates whose parents are in the file or already imported. You can include certificates in any order in the import file.
- Import the certificates one at a time in descending order. You must install the root certificate first, and then install the remaining chained certificates in descending order.

If you follow one of these methods, the system automatically chains the certificates together in the correct order and displays them hierarchically in the admin console.



If you install multiple certificates in a user's Web browser, the browser prompts the user to choose which certificate to use when signing in.

Importing Intermediate CA Certificates

To import an intermediate CA certificate:

1. Select **System > Configuration > Certificates > Device Certificates**.
2. Click the Intermediate Device CAs link to display the management page.
3. Click **Import CA certificate**.
4. Browse to the certificate file, select it, and click **Import Certificate** to complete the import operation.

Importing a Renewed Certificate That Uses the Existing Private Key

You can renew a device certificate in two ways:

- Submit a new CSR to a CA-This process is more secure because the CA generates a new certificate and private key and retires the older private key. To use this renewal method, you must first create a CSR through the admin console.
- Request renewal based on the CSR previously submitted to the CA-This process is less secure, because the CA generates a certificate that uses the existing private key.

When you order a renewed certificate, you must either resubmit your original CSR or ensure that the CA has a record of the CSR that you submitted for your current certificate.

To import a renewed device certificate that uses the existing private key:

1. Follow your CA's instructions for renewing a certificate that you previously purchased through them. Be sure to specify the same information you used in the original CSR. Your CA uses this information to create a new certificate that corresponds to the existing key.



Even though you specify the same information used in the original CSR, your root CA might have different serial numbers and keys from the original. You might need to support both new client and old client certificates during the transition period, which also requires that you maintain two root CA certificates (your existing certificate and the renewed certificate), at least temporarily

2. Select **System > Configuration > Certificates > Device Certificates**.
-

3. Click the link that corresponds to the certificate you want to renew.
4. Click **Renew Certificate** to display the page.
5. In the Renew the Certificate form, browse to the renewed certificate file, enter the password for the certificate key, and click **Import**.

Downloading a Device Certificate

You download the device certificate to your local host so that you can import it into other network devices as needed.

To download a device certificate:

1. Select **System > Configuration > Certificates > Device Certificates**.
2. Click the link of the device certificate you want to download to display the configuration page.
3. Click the **Download** link.
4. Save the file to the desired location.

Using Device Certificates with Virtual Ports

Virtual ports can be used to create multiple fully qualified domain names for user sign-in.

When a user tries to sign in using the IP address defined in a virtual port, the system uses the certificate associated with the virtual port to initiate the SSL transaction.

You can implement digital certificate security with virtual ports in either of the following ways:

- Associate all hostnames with a single certificate-With this method, you use a single wildcard certificate to validate the identity of all system hostnames, regardless of which hostname is used to sign into. A wildcard certificate includes a variable element in the domain name, making it possible for users who sign in from multiple hosts to map to the "same" domain. For example, if you create a wildcard certificate for *.yourcompany.com, the system uses the same certificate to validate its identity to users who sign in to employees.yourcompany.com as it does to users who sign into partners.yourcompany.com.

- Associate each hostname with its own certificate-With this method, you associate different hostnames with different certificates. Create a virtual port for each hostname. A virtual port activates an IP alias on a physical port. For example, you can create two virtual ports on a single appliance, mapping the first virtual port to the IP address 10.10.10.1 (sales.yourcompany.com) and the second virtual port to the IP address 10.10.10.2 (partners.yourcompany.com). Then you can associate each of these virtual ports with its own certificate, ensuring that users authenticate through different certificates.

To associate certificates with virtual ports:

1. Create the virtual ports.
2. Import the device certificates.

3. Associate the device certificates with the virtual ports:

1. Select **System > Configuration > Certificates > Device Certificates**.
2. Click the link of the device certificate you want to configure to display the configuration page.
3. Use the controls in the "Present certificate on these ports" section to associate ports with the certificate.



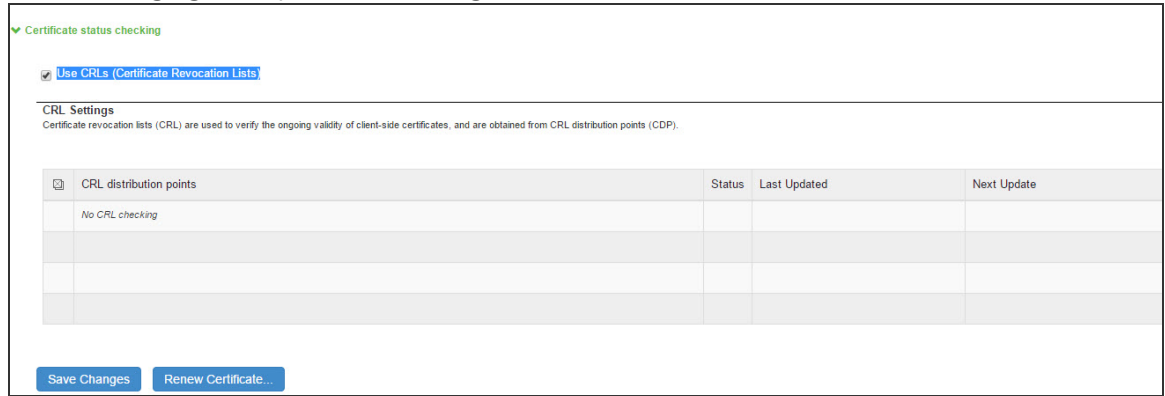
You can assign only one device certificate to the Management Port. If you assign a certificate other than the default device certificate to the Management Port, the default device certificate is automatically deselected as the default. If you do not select a device certificate for the Management Port, the system uses the default device certificate that is presented on the Internal port. You cannot assign certificates to Management Port VIPs.

Enabling Certificate Revocation Check for Device Certificate

To enable the CRL for Device Certificates:

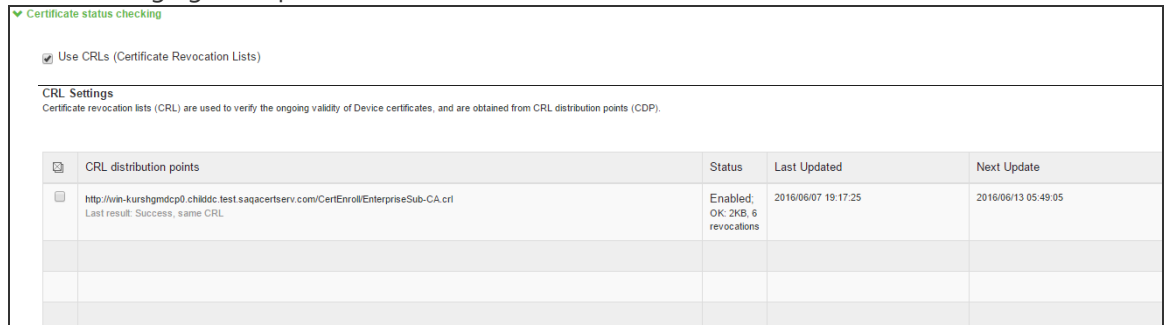
1. Go to **System > Configuration > Certificates > Device Certificates**.
2. Click on the certificate from the list to go to the certificate details.
3. In the Certificate Details page, go to Certificate Status Checking and enable the **Use CRLs (Certificate Revocation Lists)** check box.

The following figure depicts the Enabling Certificate Revocation Check for Device Certificate :



4. Click on **Save Changes**.
5. Import the CA or CA Chain that issued the Device Certificate to **System > Configuration > Trusted Server CAs**.
6. Once the CRL is successfully downloaded for Device Certificate, it is listed in the CRL distribution points.

The following figure depicts the Successful CRL Download for Device Certificate:



This version of the ICS supports the 3072 bit key length for Device Certificates.

The following figure depicts the 3072-bit Key Length for Device Certificates:

Configuration > Certificates > New Certificate Signing Request

New Certificate Signing Request

Use this page to create a new Certificate Signing Request (CSR) to send to your Certificate Authority of choice.

Common Name: <small>(e.g., secure.company.com)</small>	<input type="text" value="CN"/>
Organization Name: <small>(e.g., Company Inc.)</small>	<input type="text" value="O"/>
Org. Unit Name: <small>(e.g., IT Group)</small>	<input type="text" value="OU"/>
Locality: <small>(e.g., SomeCity)</small>	<input type="text" value="L"/>
State (fully spelled out): <small>(e.g., California)</small>	<input type="text" value="S"/>
Country (2 letter code): <small>(i.e., US)</small>	<input type="text" value="C"/>
Email Address:	<input type="text" value="E@xyz.com"/>
Key Type:	<input checked="" type="radio"/> RSA <input type="radio"/> ECC
Key Length:	<input type="text" value="1024"/> bits
	<input type="text" value="1024"/> <input type="text" value="2048"/> <input type="text" value="3072"/>
Please enter some random characters in the Random Data field using the system's random key generator. We recommend that you enter approximately twenty characters.	
Random Data: <small>(used for key generation)</small>	<input type="text"/>

Using Trusted Client CAs

This topic describes how to use trusted client Certificate Authorities (CAs).

Understanding Trusted Client CAs

A trusted client CA is a CA that you deem trusted by adding it to the trusted client CA store. The system trusts any certificate issued by that CA. To use client CA certificates, you must install and enable the proper certificates. Additionally, you must install the corresponding client-side certificates in your users' Web browsers, or you must use the MMC snap-in in your users' computer accounts (machine certificate). When validating a client-side CA certificate, the system verifies that the certificate is not expired or corrupt and that the certificate is signed by a CA that the system has been configured to recognize. If the CA certificate is chained, the system also follows the chain of issuers until it reaches the root CA, validating each issuer in turn. The system supports X.509 CA certificates in DER and PEM encode formats.

When you install a client-side certificate, you must determine whether to use the certificate to identify individual users or individual machines. To use the certificate to identify individual users, you must install the certificate in each user's individual certificate store. Then you must enable authentication using a certificate server, or you must enable authorization using realm, role, and/or resource policy settings. To use the certificate to identify individual machines, you must install the certificate in each computer's certificate store. Then you must configure a Host Checker policy that checks for the machine certificate and authorizes access to realms, roles, or resource policies based on the certificate's validity.

The system supports using the following additional features with CA certificates:

- **Certificate servers**-A certificate server is a type of local authentication server that allows you to authenticate users based solely on their certificate attributes rather than authenticating them against a standard authentication server (such as LDAP or RADIUS), and it requires specific certificates or certificate attributes.
- **Certificate hierarchies**-Within a certificate hierarchy, one or more subordinate certificates (called intermediate certificates) are branched off a root certificate to create a certificate chain. Each intermediate certificate (also called a chained certificate) handles requests for a part of the root CA domain. For example, you can create a root certificate that handles all requests to the `yourcompany.com` domain and then branch off intermediate certificates that handle requests to `partners.yourcompany.com` and `employees.yourcompany.com`. When you install a chained certificate, the system confirms that the chain is valid and allows users to authenticate using the leaf certificate (that is, the lowest certificate in the chain).
- **Certificate revocation lists**-Certificate revocation is a mechanism by which a CA invalidates a certificate before its expiration date. The CA publishes a certificate revocation list (CRL) which is a list of revoked certificates. Within CRLs, each entry contains the serial number of the revoked certificate, the date that the certificate was revoked, and the reason the certificate was revoked. The CA can invalidate a certificate for various reasons such as when the employee to whom the certificate is issued leaves the company, the certificate's private key is compromised, or the client-side certificate is lost or stolen. When the CA revokes a certificate, the system can appropriately deny access to users who present a revoked certificate.

Trusted Client CA Implementation Notes

Uploading a trusted client CA certificate does not enable client-side SSL authentication or authorization. To do so, you must use a certificate server, or enable certificate restrictions at the realm, role, or resource policy level, or create a Host Checker policy that verifies a machine certificate.

With client-side certificates, we strongly recommend that you advise users to close their Web browsers after signing out. If they do not, other users might be able to use their open browser sessions to access certificate-protected resources without reauthentication. After loading a client-side certificate, Internet Explorer caches the certificate's credentials and private key. The browser keeps this information cached until the user closes the browser (or, in some cases, until the user reboots the workstation). For details, see <http://support.microsoft.com/?kbid=290345>. To remind users to close their browsers, you can modify the sign out message on the Sign-in Pages tab.

Certificate authentication does not work on Internet Explorer 8, 9, and 11 if SSL 2.0 is enabled with other SSL and TLS versions. For details, see <http://support.microsoft.com/kb/2851628>.

Understanding CRLs

A certificate revocation list (CRL) is a mechanism for canceling a client-side certificate. As the name implies, a CRL is a list of revoked certificates published by a CA or a delegated CRL issuer. The system supports base CRLs, which include all of the company's revoked certificates in a single, unified list.

The system determines the correct CRL to use by checking the client's certificate. (When it issues a certificate, the CA includes CRL information for the certificate in the certificate itself.) To ensure that it receives the most up-to-date CRL information, the system periodically contacts a CRL distribution point to get an updated list of CRLs. A CRL distribution point (CDP) is a location on an LDAP directory server or Web server where a CA publishes CRLs. The system downloads CRL information from the CDP at the interval specified in the CRL, at the interval that you specify during CRL configuration, and when you manually download the CRL. The system also supports CRL partitioning. CRL partitioning enables you to verify portions of very large CRLs without spending the time and bandwidth necessary to access and validate a very large CRL or collection of large CRLs. CRL partitioning is only enabled when you employ the Specify the CDP(s) in the client certificates method (described below). In this case, the system validates the user by verifying only the CRL specified in the client certificate.

Although CAs include CRL information in client-side certificates, they do not always include CDP information as well. A CA can use any of the following methods to notify the system of a certificate's CDP location:

- Specify the CDP(s) in the CA certificate-When the CA issues a CA certificate, it might include an attribute specifying the location of the CDPs that the system should contact. If more than one CDP is specified, the system chooses the first one listed in the certificate and then fails over to subsequent CDPs, if necessary.

- Specify the CDP(s) in the client certificates-When the CA issues a client-side certificate, it might include an attribute specifying the location of the CDPs that the system must contact. If more than one CDP is specified, it chooses the first one listed in the certificate and then fails over to subsequent CDPs, if necessary. When the system employs CRL partitioning and the client certificate specifies only one CRL, it performs verification using only that CRL.



If you choose this method, the user receives an error on the first sign-in attempt because no CRL information is available. Once the system recognizes the client's certificate and extracts the CRL location, it can start downloading the CRL and subsequently validate the user's certificate. In order to successfully sign in, the user must try to reconnect after a few seconds.

- Require the administrator to manually enter the CDP location-If the CA does not include the CDP location in the client or CA certificates, you must manually specify how to download the entire CRL object. You can specify a primary and backup CDP. (Manually entering the CDP location provides the greatest flexibility because you do not need to reissue certificates if you change the CDP location.)

The system compares the user's certificate against the appropriate CRL during authentication. If it determines that the user's certificate is valid, the system caches the certificate attributes and applies them, if necessary, during role and resource policy checks. If it determines that the user's certificate is invalid, if it cannot contact the appropriate CRL, or if the CRL is expired, it denies the user access.



- The system supports only CRLs that are in a PEM or DER format and that are signed by the CA for which the revocations apply.
 - The system only saves the first CRL in a PEM file.
-

Understanding OCSP

The Online Certification Status Protocol (OCSP) is a service that enables you to verify client certificates. When OCSP is enabled, the system becomes a client of an OCSP responder and forwards validation requests for users based on client certificates. The OCSP responder maintains a store of CA-published certificate revocation lists (CRLs) and maintains an up-to-date list of valid and invalid client certificates. After the OCSP responder receives a validation request, it validates the status of the certificate using its own authentication database, or it calls upon the OCSP responder that originally issued the certificate to validate the request. After formulating a response, the OCSP responder returns the signed response, and the original certificate is either approved or rejected.

Importing a Trusted Client CA Certificate

If you require users to provide a client-side certificate to sign in, you must upload the corresponding CA certificate. You can upload CA certificates manually, or you can configure the system to upload CA certificates automatically. The system uses the uploaded certificate to verify that the browser-submitted certificate is valid. In addition, you can specify whether or not to automatically import CA certificates for validation, and you can specify a CRL or OCSP retrieval method to use to automatically import CA certificates.

To import a trusted client CA certificate:

1. Select **System > Configuration > Certificates > Trusted Client CAs** to display the configuration page.
2. Click **Import CA Certificate** to display the configuration page.
3. Browse to the certificate file, select it, and click **Import Certificate** to complete the import operation.

Renewing a Certificate

To renew a certificate:

1. Select **System > Configuration > Certificates > Trusted Client CAs**.
2. Click the link for the certificate you want to renew.
3. Click **Renew Certificate** to display the import certificate page.
4. Browse to the certificate file, select it, and click **Import Certificate** to complete the import operation.

Configuring Auto-Importing of Client Certificates

To enable auto-importing:

1. Select **System > Configuration > Certificates > Trusted Client CAs**.
2. Click the **Auto-Import Options** button to display the options.
3. Complete the configuration described in the following table.
4. Save your changes.

The following table lists the Auto-Import Options Settings:

Settings	Guidelines
Auto-import trusted CAs	Select this option to enable auto-import and display its configuration settings.
Client Certificate Status Checking	Select a method to validate the trusted client certificate: None-Do not validate. Use OCSP-Use the OCSP method, validating the client certificate in real-time, as needed. After you select this option, you can specify options for OCSP. Use CRLs-Use CRLs to validate the client certificate. After you select this option, you can specify options for CRL. Use OCSP with CRL fallback-Use the OCSP validation method when possible, but attempt to validate client certificates using CRLs if the OCSP method fails (for example, if the link to the OCSP responder fails). After you select this option, you can specify options for OCSP and CRL. Inherit from root CA-Use the method configured for the device certificate.
CDP(s)/OCSP responder	Select the location of the responder value: None-Do not use the responder. From client certificate-Use the responder value configured in the client certificate. From trusted CA certificate-Use the responder value configured in the trusted CA certificate that has been uploaded to the system.
Verify imported CA certificates	Select this option to verify that this trusted client CA is valid. Enabling this will check the CRL of this certificate's issuer, and repeat up the chain until reaching the root trusted client CA.

Settings	Guidelines
Skip Revocation check when OCSP/CDP server is not available	<p>Select this option to instruct ICS to skip revocation check and accept end user certificates when either OCSP server or CDP server is not accessible over the network. This option is applicable to digital certificates used for end user authentication.</p> <p>The OCSP Timeout, applicable only for OCSP, is used as the maximum timeout for a network connection or data transfer operation while connecting to an OCSP Responder. An internal timeout will be used for CDP.</p> <p>ICS skips the revocation check in the following conditions:</p> <ul style="list-style-type: none"> Server IP is not reachable Server Hostname is either not resolvable or resolving to non OCSP/CRL Server IP Proxy IP is either not reachable or not resolving Downloaded CRL has expired OCSP/CRL service in Server is not responding

Configuring Options for Trusted Client CA Certificates

To configure options for the trusted client CA certificate:

1. Select **System > Configuration > Certificates > Trusted Client CAs**.
2. Click the certificate you want to configure.
3. Complete the configuration described in the following table.

The following table lists the Trusted Client CA Settings:

Settings	Guidelines
Certificate	<p>Use the expander buttons to display the following details:</p> <p>Issued To-Name and attributes of the entity to whom the certificate is issued.</p> <p>Issued By-Name and attributes of the entity that issued the certificate. Note that the value of this field must match either the Issued To field (for root certificates) or the Issued To field of the next highest certificate in the chain (for intermediate certificates).</p> <p>Valid Dates-Time range for which the certificate is valid.</p> <p>Details-Variou certificate details, including its version, serial number, signature algorithm, CRL distribution points, public key algorithm type, and public key.</p>
Client Certificate Status Checking	<p>Select a method to validate the trusted client certificate:</p> <p>None-Do not validate.</p> <p>Use OCSP-Use the OCSP method, validating the client certificate in real-time, as needed. After you have selected this option and saved the configuration, you can specify options for OCSP.</p> <p>Use CRLs-Use CRLs to validate the client certificate. After you have selected this option and saved the configuration, you can specify options for CRL.</p> <p>Use OCSP with CRL fallback-Use the OCSP validation method when possible, but attempt to validate client certificates using CRLs if the OCSP method fails (for example, if the link to the OCSP responder fails). After you have selected this option and saved the configuration, can specify options for OCSP and CRL.</p> <p>Inherit from root CA- Use the method configured in Root CA. This Option is only applicable for Intermediate CA.</p>
Verify Trusted Client CA	<p>Select this option to verify that this trusted client CA is valid. Enabling this will check the CRL of this certificate's issuer, and repeat up the chain until reaching the root trusted client CA.</p>
Trusted for Client Authentication	<p>Clear this check box to exclude the CA from being trusted for client certificate authentication. You might want to do this if this CA was added for another trusting purpose, such as SAML signature verification or machine certificate validation.</p>

Settings	Guidelines
Participate in Client Certificate Negotiation	<p>This option is available only on Ivanti Connect Secure.</p> <p>Select this option to include the CA participation in client certificate selection for authentication.</p> <p>In client certificate authentication or restriction, the device sends a list of all trusted client CAs configured in the trusted client CA store with this flag enabled to the user's browser for user certificate selection. The browser prompts the client certificates whose issuer CA and/or root CA is in that list. This option allows you to control which client certificate(s) are prompted for selection. Clearing this option for all certificates in a CA chain results in those certificates not being prompted.</p>
Skip Revocation check when OCSP/CDP server is not available	<p>Select this option to instruct ICS to skip revocation check and accept end user certificates when either OCSP server or CDP server is not accessible over the network. This option is applicable to digital certificates used for end user authentication.</p> <p>The OCSP Timeout, applicable only for OCSP, is used as the maximum timeout for a network connection or data transfer operation while connecting to an OCSP Responder. An internal timeout will be used for CDP.</p> <p>ICS skips the revocation check in the following conditions:</p> <ul style="list-style-type: none"> Server IP is not reachable Server Hostname is either not resolvable or resolving to non OCSP/CRL Server IP Proxy IP is either not reachable or not resolving Downloaded CRL has expired OCSP/CRL service in Server is not responding

4. Save your changes.
5. If you have enabled **CRL Checking**, click **CRL Checking Options**.
6. If you have enabled OCSP options:
 1. Click **OCSP Options**.
 2. Complete the configuration described in following table.

The following table lists the OCSP Options Settings:

Settings	Guidelines
Use	<p>Select the type of OCSP responder to validate trusted client CAs:</p> <p>None-The system does not use OCSP to verify the status of certificates issued by this CA.</p> <p>Responder(s) specified in the CA certificate-The system uses OCSP responders specified in the imported client CA to perform verification. When you select this option, the system displays a list of OCSP responders specified in the imported CA (if any) and the last time they were used.</p> <p>Responder(s) specified in the client certificates-The system uses responders specified during client authentication to perform verification. When you select this option, the system displays a list of known OCSP responders (if any) and the last time they were used.</p> <p>Manually configured responders-The system uses primary and secondary OCSP responders at the addresses you specify.</p>
Device Certificate to sign the request	Select the appropriate device certificate or leave the default (unsigned).
Signature Hash Algorithm	Select SHA-1 or SHA-2.
Use Nonce	A nonce is random data the system includes in an OCSP request and the OCSP responder returns in the OCSP response. The system compares the nonce in the request and response to ensure that the response is generated by the OCSP responder. If the two do not match, the system disregards the response and sends a new request. Nonces are a common way of preventing replay attacks.

7. Save the configuration.
8. After you have added an OCSP responder to the list, you can click its link to display the page.
9. Complete the configuration described in Table 127.

The following table lists the Responder Signer Certificate Settings:

Settings	Guidelines
Responder Signer Certificate	Browse to the network path or local directory location of a Responder Signer Certificate. This is the certificate the OCSP responder uses to sign the response. You must specify the Responder Signer Certificate if the signer certificate is not included in the response.
Trust Responder Certificate	Select this option to allow an OCSP responder certificate that matches the responder signer certificate.
Revocation Checking	Select this option to ensure that the certificate has not recently been revoked. This option has implications only if you specified the Use OCSP with CRL fallback option.
Allow clock discrepancy	Use this option to account for possible mismatches in timestamps between the system clock and the OCSP responder clock. If the mismatch is significant, the system disregards the response from the OCSP responder as out of date or expired.

10. Save the configuration.

Configuring a Proxy Server for CRL Downloads and OCSP Status Checks

You can configure the system to send CRL download requests and OCSP status checks to the proxy server and collect the response. You might want to do this if you deploy proxy server to control access to the Internet.

The following types of CRL downloads can use the proxy server:

- CRL distribution points (CDPs) specified in the trusted client CAs
- CDPs specified in client certificates
- Manually configured CDPs

Similarly, the system can send OCSP requests to the OCSP responder through the proxy server. The OCSP responses are also received through the proxy server. This feature is useful when you deploy a large number of Ivanti access systems and the OCSP responders are located outside the network.

To configure a proxy server:

1. Select **System > Configuration > Certificates > Trusted Client CAs**.

2. Click **Proxy Settings** to display the page.
3. Complete the configuration described in Table 128.
4. Save the configuration.

The following table lists the Proxy Settings:

Settings	Guidelines
Use Proxy Server for HTTP-based CRL download	Select to enable the CRL operations to use a proxy server. You can configure a proxy server for web-based URLs, not LDAP URLs.
Use Proxy Server for HTTP-based OCSP status checking	Select to enable the OCSP operations to use a proxy server.
Host Address	Specify either an IP address or a fully qualified domain name.
Port	Enter the proxy server port number if it is different from the default value of 80.
Username/password	If your proxy server required authentication, enter a username and password to log in to the proxy server.

Using Trusted Server CAs

This topic describes trusted server certificate authorities (CAs).

Understanding Trusted Server CAs

All of the trusted root CAs for the Web certificates installed in Internet Explorer are preinstalled. You might need to install a trusted server CA for additional Web servers in the following situations:

- If you are using third-party integrity measurement verifiers (IMVs) that are installed on a remote server, you must upload the trusted root certificate of the CA that signed the remote server's server certificate.
- If you are using virus signature version monitoring with your own staging site for storing the current virus signatures list, you must upload the trusted root certificate of the CA that signed the staging server certificate.

You can install the trusted root CA certificate on the endpoint in any of the following ways:

- Use a CA certificate that is chained to a root certificate that is already installed on the endpoint, such as VeriSign.
- Upload the CA certificate and any intermediate CA certificates to the Pulse client system. During client installation, the system automatically installs the trusted root device CA certificates on the endpoint. When prompted during installation, the user must allow the installation of the CA certificate(s).
- Prompt users to import the CA certificates on the endpoint using Internet Explorer or other Microsoft Windows tools. In other words, you can use common methods organizations use to distribute root certificates.



You cannot use CRL revocation checks for trusted server CA certificates.

Uploading Trusted Server CA Certificates

You can use the Trusted Server CAs page to upload the trusted root certificate of the CA that signed the Pulse client service device certificate. If you upload a certificate chain, you must install the certificates one at a time in descending order starting with the root certificate (DER or PEM files), or you must upload a single file that contains the entire certificate chain (PEM files only). The system supports X.509 CA certificates in PEM (Base 64) and DER (binary) encode formats.

To upload CA certificates:

1. Select **System > Configuration > Certificates > Trusted Server CAs** to display the page.
2. Click **Import Trusted Server CA** to display the page.
3. Browse to the certificate file, select it, and click **Import Certificate** to complete the import operation.

Restoring the Prepopulated Group of Trusted Server CA Certificates

The System > Configuration > Certificates > Trusted Server CAs page is prepopulated with some of the trusted root CAs for the Web certificates installed in Internet Explorer and Windows. You can use the delete functionality on this page to delete CAs and the reset functionality to restore the list to the set that was installed during the upgrade. The reset operation clears all manually imported certificates.

To restore the prepopulated group of trusted CA certificates:

1. Select **System > Configuration > Certificates > Trusted Server CAs**.
2. Click **Reset Trusted Server CAs**.

3. Confirm that you want to restore the set of trusted server CAs that was installed when you upgraded.

Ivanti Connect Secure restores the group of prepopulated trusted server CAs that were installed upon upgrade. This operation clears all manually imported certificates.

Renewing a Trusted Server CA Certificate

If a trusted CA renews its certificate, you must upload the renewed CA certificate.

To import a renewed CA certificate:

1. Select **System > Configuration > Certificates > Trusted Server CAs**.
2. Click the link that corresponds to the certificate that you want to renew to display the page.
3. Click **Renew Certificate**.
4. Browse to the certificate file, select it, and click **Import Certificate** to complete the import operation.

Deleting a Trusted Server CA Certificate

You can delete any trusted server CA certificate, including preinstalled certificates.

To delete a trusted server CA certificate:

1. Select **System > Configuration > Certificates > Trusted Server CAs**.
2. Select the check box for the certificate you want to delete.
3. Click **Delete**, and then confirm that you want to delete the certificate.

Using Code-Signing CAs

- This topic describes how to use code-signing Certificate Authorities (CAs).

Understanding Code-Signing CAs

In a basic setup, the only required certificates are a device certificate and a code-signing certificate. Ivanti Connect Secure can use a single code-signing certificate to resign all Java applets and a single device certificate to intermediate all other PKI-based interactions. If the basic certificates do not meet your needs, however, you may install multiple device and applet certificates on Ivanti Connect Secure or use trusted CA certificates to validate users.

When Ivanti Connect Secure intermediates a signed Java applet, it re-signs the applet with a self-signed certificate by default. This certificate is issued by a nonstandard trusted root CA. As a result, if a user requests a potentially high-risk applet (such as an applet that accesses network servers), the user's Web browser alerts him that the root is untrusted.

If you import a code-signing certificate, Ivanti Connect Secure uses the imported certificate to re-sign applets instead of the default self-signed certificate. As a result, if a user requests a potentially high-risk applet, the user's Web browser displays an informational message instead of a warning. The message informs the user that the applet is signed by a trusted authority.

The system supports the following types of code-signing certificates:

- **Microsoft Authenticode Certificate**-The system uses this certificate to sign applets that run on either Microsoft JVM or Oracle JVM. Note that we only support Microsoft Authenticode Certificates issued by Verisign.
- **JavaSoft Certificate**-The system uses this certificate to sign applets that run on Oracle JVM. Note that we only support JavaSoft Certificates issued by Verisign and Thawte.

When deciding which code-signing certificate to import, consider the following browser dependencies:

- **Internet Explorer**-Internet Explorer running on new computers shipped with Windows pre-installed typically runs the Oracle JVM, which means that Ivanti Connect Secure needs to re-sign applets using the JavaSoft certificate.

Internet Explorer running on an older version of Windows that has been upgraded may run the Microsoft JVM, which means that Ivanti Connect Secure needs to re-sign applets using the Authenticode certificate.
- **Netscape, Firefox, and Safari**-These browsers only support the Oracle JVM, which means that Ivanti Connect Secure needs to re-sign applets using the JavaSoft certificate.

Additional Considerations for Oracle JVM Users

By default, the Java Plug-in caches an applet along with the code-signing certificate presented when a user accesses the applet. This behavior means that even after importing a code-signing certificate to Ivanti Connect Secure, the browser continues to present applets with the original certificate. To ensure that JVM users are not prompted with an untrusted certificate for applets accessed prior to importing a code-signing certificate, users need to flush the Java Plug-in cache. Alternatively, users can disable the cache, but this option may impact performance since the applet needs to be fetched each time the user accesses it.

The Java Plug-in maintains its own list of trusted Web server certificates that is different from the browser's list of trusted certificates. When a user accesses an applet, the JVM makes its own connection (in addition to the browser) to the Web server on which the applet resides. The user is then presented with the option to accept the Web server certificate in addition to the code-signing certificate. In these cases, the user needs to select the Always Trust button for the Web server certificate. Due to a built-in timeout in the Java Plug-in, if the user waits too long to select this button for the Web server certificate, the applet does not load.

Importing a Code-Signing CA Certificate

To import a code-signing certificate:

1. Select **System > Configuration > Certificates > Code-Signing Certificates** to display the configuration page.
2. Click **Import Certificates** to display the configuration page.
3. Complete the configuration described in the following table.

The following table lists the Import Certificates Configuration Guidelines:

Settings	Guidelines
Microsoft Authenticode or Multipurpose Certificate for Internet Explorer (Microsoft JVM)	
Certificate File	Browse to the network path or local directory location of your certificate key file.
Private Key File	Browse to the network path or local directory location of your private key file.
Password Key	Enter the password key.
Javasoftware Certificate for Internet Explorer & Netscape (Sun JVM)	

Settings	Guidelines
Keystore File	Browse to the network path or local directory location of the keystore file.
Password key	Enter the password key.

4. Click **Import** to complete the import operation.
5. When you have successfully imported a certificate, the system displays the Sign Ivanti Web Controls With dialog box. Specify the signing option:
 - Default Ivanti Certificate-Select this option to sign all ActiveX and Java applets originating from Ivanti Connect Secure using the default Ivanti certificate. If you have previously selected an imported code-signing certificate and are reverting back to this option, after you click Save, a process icon appears indicating that the system is processing the request and re-signing all of the relevant code. This process can take several minutes to complete.
 - Authenticode Certificate- For <Imported Certificate Name> -Select this option to sign all ActiveX and Java applets using the certificate or certificates imported in the previous step. When you click Save, a process icon appears indicating that the system is processing the request and signing all of the relevant code. This process can take several minutes to complete.
6. Use settings in the following tabs to specify which resources are signed or re-signed by the applet certificate:
 - Users > Resource Policies > Web > Java > Code Signing

Using Code-Signing Certificates for Java Applets

To use code-signing certificates with Java applets:

1. Install the Java code-signing certificates. Use the System > Configuration > Certificates > Code-Signing Certificates page.
2. Use any of the following methods:
 - Create code-signing policies specifying which applets to re-sign. Use the Users > Resource Policies > Web > Java > Code Signing page or the Users > Resource Profiles > Web Application Resource Profiles > Profile page. The policies must specify the hostnames from which the applets originate.

- Upload your own Java applets to Ivanti Connect Secure and configure the system to sign or re-sign the applets.

Using Client Auth Certificates

This topic describes how to use client auth certificates.

Understanding Client Auth Certificates

In certain corporate environments, servers on the LAN are protected with two-way SSL authentication. These servers require the client to authenticate by presenting a valid certificate.

In the remote access scenario, Ivanti Connect Secure is a client of these servers. You can configure Ivanti Connect Secure to present client authentication certificates to servers whenever it communicates over SSL. Note that Ivanti Connect Secure will present client certificates only when the SSL handshake requires it.



This feature authenticates Ivanti Connect Secure (as a client) to back-end servers. It also authenticates end users or end-user machines to servers on the corporate LAN.

The SSL protocol provides for mutual authentication of server and client at the time of session initiation. The client part of the authentication is optional. For enhanced security, some deployments may require that the client also authenticate itself with a certificate. Normally, when setting up an SSL connection with a server on behalf of the end user, Ivanti Connect Secure does not present any certificate to the server. It needs to be explicitly configured to present such certificate. This section explains how such configuration may be performed.

The basic idea is to upload a certificate, private key pair to the Ivanti access management framework and configure a mapping between this pair and a server resource. Subsequently, when an end user attempts to establish a connection with that server, Ivanti Connect Secure presents the associated certificate to the server. If no certificate is associated with the server in Ivanti Connect Secure's certificate store, then it is assumed that the server does not demand client certificate.

If, during the SSL handshake, the back-end server requests a client certificate but Ivanti Connect Secure doesn't send a certificate, the end user sees an "access denied" error message. Similarly, if the back-end server rejects the Ivanti Connect Secure certificate, the end user sees an "access denied" error message. If a certificate is configured, is successfully retrieved and no error is encountered during handshake, the user is granted access to the server.



The Ivanti access management framework allows client authentication certificates to be uploaded to the device in two ways: generate a CSR and upload the signed certificate returned by the CA, or directly import the certificate if one is available.

Importing a Client Auth Certificate

The Ivanti access management framework allows certificates that include the private key and for instances where the private key is in a separate file from the certificate. In addition, if your certificates have been exported into a system configuration file, you can import the system configuration file to upload the certificates.

To import the client auth certificates files:

1. Select **System > Configuration > Certificates > Client Auth Certificates**.
2. Click **Import Certificate & Key** to display the configuration page.
3. Complete the configuration described in Table 130.
4. Click **Import**.

The following table lists the Import Certificate and Key Settings:

Settings	Guidelines
If certificate file includes private key	
Certificate File	Browse to the network path or local directory location of your private key file.
Password Key	Enter the password key.
If certificate and private file are separate keys	
Certificate File	Browse to the network path or local directory location of your certificate key file.
Private Key File	Browse to the network path or local directory location of your private key file.
Password Key	Enter the password key.
Import via System Configuration file	

Settings	Guidelines
System Configuration File	Browse to the network path or local directory location of the system configuration file.
Password	Enter the password.

Renewing a Client Auth Certificate

To renew a certificate:

1. Select **System > Configuration > Certificates > Client Auth Certificates**.
2. Click the link that corresponds to the certificate you want to renew.
3. Click **Renew Certificate** to display the configuration page.
4. In the Renew the Certificate form, browse to the renewed certificate file, enter the password for the certificate key, and click **Import**.

Configuring Two-Way SSL Authentication

To configure two-way SSL authentication:

1. Import the certificates used for two-way SSL handshake in the **System > Configuration > Certificates > Client Auth Certificates** window.
2. Define the back-end resource and assign a certificate to be presented when accessing it using the **Users > Resource Policies > Web > Client Authentication** window.

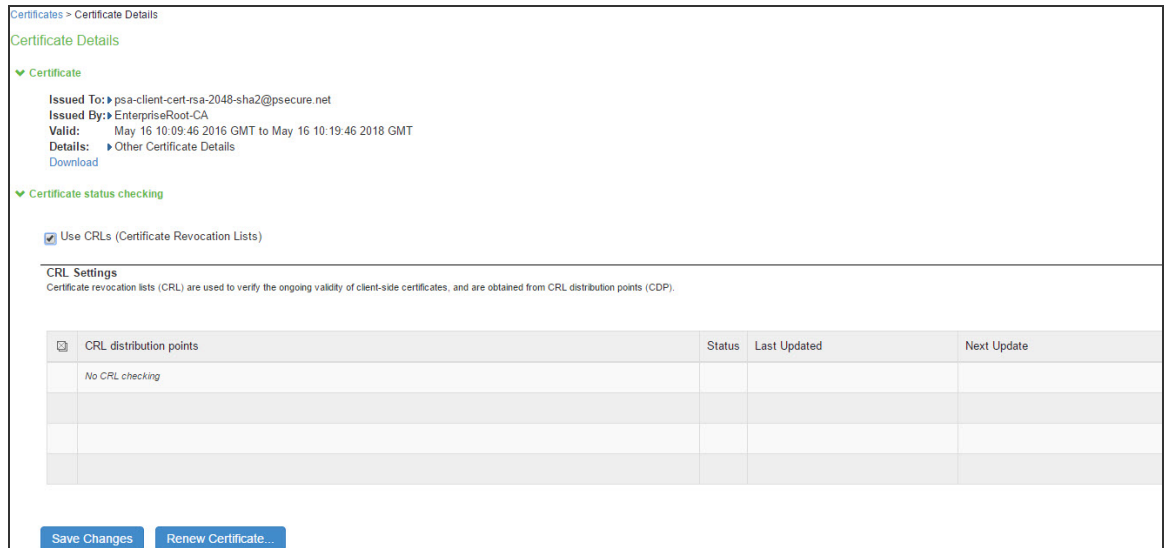
Enabling Certificate Revocation Check for Client Auth Certificate

Client Auth Certificate Revocation Check is only applicable for TLS Syslog Backend Server. It is not applicable for any other backend server configured to ask Client Certificate.

To enable the CRL for Client Auth Certificate:

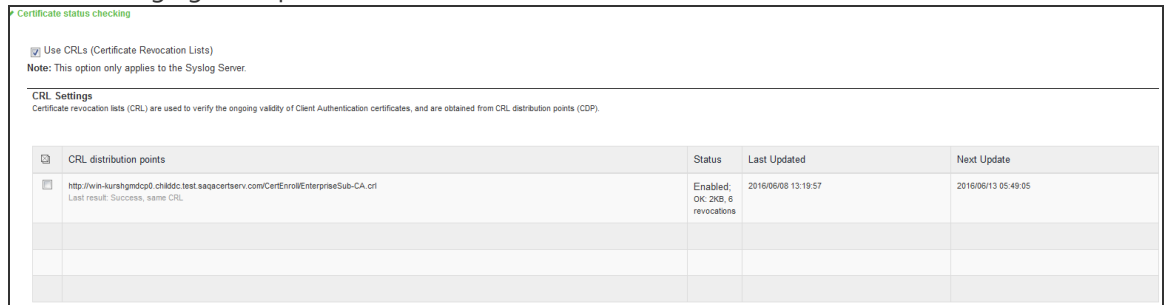
1. Go to **System > Configuration > Certificates > Client Auth Certificates**.
2. Click on the certificate from the list to go to the certificate details.
3. In the Certificate Details page, go to **Certificate Status Checking** and enable the **Use CRLs (Certificate Revocation Lists)** check box.

The following figure depicts Enabling Certificate Revocation Check for Client Auth Certificate:



4. Click on **Save Changes**.
5. Import the CA or CA Chain that issued the Client Auth Certificate to **System -> Configuration -> Trusted Client CAs**.
6. Once the CRL is successfully downloaded for Client Auth Certificate, it is listed in the CRL distribution points. See Figure.

The following figure depicts Successful CRL Download for Client Auth Certificate:



This version of the ICS supports the 3072-bit key length for Client Auth Certificates. See following figure.

The following figure depicts the 3072-bit Key Length for Client Auth Certificates:

Configuration > Certificates > New Certificate Signing Request

New Certificate Signing Request

Use this page to create a new Certificate Signing Request (CSR) to send to your Certificate Authority of choice

Common Name:
(e.g., secure.company.com)

Organization Name:
(e.g., Company Inc.)

Org. Unit Name:
(e.g., IT Group)

Locality:
(e.g., SomeCity)

State (fully spelled out):
(e.g., California)

Country (2 letter code):
(i.e., US)

Email Address:

Key Type: RSA ECC

Key Length: bits

Please enter some random characters in the **3072** system's random key generator. We recommend that you enter approximately twenty characters.

Random Data:
(used for key generation)



CRL Download for Device Certificate and Client Auth Certificate using LDAP based URL won't work due to dependency of LDAP Username and Password. In some cases, CDP LDAP URL hostname field is also required which is also not supported.

Mapping Resource Policies to the Certificate

Once the certificates have been uploaded, you can map resources to the certificates and the roles to which they apply.

1. Select Users > Resource Policies > Web > Client Authentication.
2. **If you do not see the Client Authentication menu item, select Users > Resource Policies > Web.**
 1. Click the Customize button in the upper right corner of the console.
 2. In the Customize View dialog box, select Client Authentication.
 3. Click OK.

4. Click the Client Authentication tab.
5. Click New Policy.
6. On the New Policy page:
 - Enter a name to label this source interface policy.
 - Enter an optional description.
7. In the Resources section, specify the back-end servers to which this policy applies. Valid values/formats are: hostnames, IP addresses, IP Address:Port and Hostname:Port.

If you specify * as the resource, one certificate is used for all resources requesting a back-end certificate authentication. This certificate becomes the default certificate. Defining a default certificate is not required.
8. In the Roles section, select one of the following options:
 - Policy applies to ALL roles-To apply this policy to all users.
 - Policy applies to SELECTED roles-To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
 - Policy applies to all roles OTHER THAN those selected below-To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
9. In the Action section, select one of the following options:
 - Use the Client Authentication Certificate Below-Select this option to associate this policy with a client authentication certificate. Select the certificate to use from the Certificate menu.
 - If the Certificates menu is blank, no certificates have been uploaded to the System > Configuration > Certificates > Client Auth Certificates window.
 - Do not use Client Authentication-If this option is selected, the system does not perform client authentication for the configured resource.
 - Use Detailed Rules-Select this option to specify one or more detailed rules for this policy.
10. Click Save Changes.

Mapping a Client Authentication Auto-Policy

A client authentication auto-policy option is available on the Users > Resource Profiles > Web page. If the back-end server requires two-way SSL authentication, this auto-policy lets you configure a certificate to be presented during the SSL handshake.

1. Select Users > Resource Profiles > Web.
2. Follow the process as a regular resource profile for defining the name and type.
3. Select the Autopolicy: Client Authentication check box.
4. In the Resource field, specify the back-end server. Valid formats/values are: hostnames, IP addresses, IP Address:Port, and HostName: Port.

If you specify * as the resource, one certificate is used for all resources requesting a back-end certificate authentication. This certificate becomes the default certificate. Defining a default certificate is not required.

5. Click Save Changes.

Checking Certificate Expiry

Every time a certificate is added to ICS (through manual import, XML import, or upgrade), its expiration date is stored in the cache. A background process checks all certification expiration dates once in every 7 days. If any certificate is about to expire soon, the administrator is notified. Notifications to administrators include a banner message in the adminUI upon login, SNMP trap, and log messages in the event log. The administrator can configure how soon he or she wishes to be notified of the expiration. The default is 60 days in advance. It can be configured to a value starting from 7 days in advance to 999 days in advance of the expiration of the certificate. The expiration warning window is common to all types of certificates. However, the administrator can choose to enable or disable this feature for each certificate category in the user interface.

Features of Certificate Expiry Warning

An administrator can know about certificates that are going to expire in the near future and avoid any unexpected downtime due to certificate expiry. Administrators can take corrective actions whenever a certificate is about to expire in order to ensure there is no service disruption.

- An administrator can enable/disable this feature for each category of certificates.

- An administrator can set how many days in advance I should be notified about certificate expiry. This is common to all certificate types.
- Read only administrators are not allowed to change these values.
- This feature is enabled by default just after upgrade or a Binary/XML import.
- When an administrator logs in to a cluster, the certificate expiration warning messages is seen for both nodes.

To check validity of certificates:

1. Click on Configuration-> Certificates -> Certificates Validity Check.
2. The page displays the Certificate Expiry Notification Duration and the Certificate Types.
3. Enter the number of days before which the warning must be displayed.
4. Select the type of certificate for which the expiry notification is required. By default, all types of certificates will be selected if no selection is made.

The following figure shows certificates validity check page.:

Configuration > Certificates > Certificates Validity Check

Certificates Validity Check

Configuration

Licensing | Pulse One | Security | **Certificates** | DMI Agent | NCP | Sensors | Client Types | Pulse Collaboration | Virtual Desktops

User Record Synchronization | IKEV2 | SAML | Mobile | VPN Tunneling

Device Certificates | Trusted Client CAs | Trusted Server CAs | Code-signing Certificates | Client Auth Certificates | **Certificates Validity Check**

Certificate Expiry Notification Duration
The certificates that have already expired or will expire within the next 60 days will be displayed by default. This time boundary can be changed by updating the value in the Certificate Expiry Notification Duration.


Certificate Expiry Notification Duration: Days 7-999 Valid range of Number of Days

Certificate Types
Select the type of certificates for which the expiry notification is required. By default, all types of certificates will be selected if no selection is made.

Device Certificates
 Trusted Client CAs
 Trusted Server CAs
 Code-Signing Certificates
 Client Auth Certificates


5. Click on Check Now. The Certificate Category, DN name and date of expiry are displayed as seen in the following figure.

The following figure depicts the Certificate Expiration Page:



- When an administrator logs in, a warning sign is displayed, if there are any certificates that expire within the configured number of days.

The following figure depicts the Warning Signal Displayed (**System > Status > System Status Overview**):



- To check if the certificate expiry warning is logged, click on log monitoring. The certificate expiry warning logs are displayed.

The following figure depicts the Certificate Expiry Warning Logs (On Main Menu bar: **System**):

Info	STS30667	2016-05-01 22:00:28 - ive - [127.0.0.1] System()[] - Number of NCP connections: 0
Info	STS20642	2016-05-01 22:00:28 - ive - [127.0.0.1] System()[] - Number of concurrent mail users logged in to the email proxy: 0
Info	STS20641	2016-05-01 22:00:28 - ive - [127.0.0.1] System()[] - Number of concurrent users logged in to the device: 0
Major	SYS31211	2016-05-01 21:14:35 - ive - [127.0.0.1] System()[] - Certificate 'CA Disig' is about to expire within next 60 days
Major	SYS31211	2016-05-01 21:14:35 - ive - [127.0.0.1] System()[] - Certificate 'Swisskey Root CA' is about to expire within next 60 days
Info	STS30667	2016-05-01 21:00:07 - ive - [127.0.0.1] System()[] - Number of NCP connections: 0
Info	STS20642	2016-05-01 21:00:06 - ive - [127.0.0.1] System()[] - Number of concurrent mail users logged in to the email proxy: 0
Info	STS20641	2016-05-01 21:00:06 - ive - [127.0.0.1] System()[] - Number of concurrent users logged in to the device: 0

- Already expired certificates under the tabs Device Certificates, Trusted client CAs, Trusted Server CAs and Client auth certificates are displayed in red color.
- For code signing certificates, if it has expired, a string "EXPIRED" is displayed in red color. The image below displays code signing certificates that have expired.

Elliptic Curve Cryptography

Understanding ECC Certificates

Public-key cryptography is a cryptographic system that requires a secret key and a public key that are mathematically linked with each other. One key encrypts the plain text while the other decrypts the cipher text. RSA is the most widely used public-key algorithm.

Elliptic Curve Cryptography (ECC) were introduced as an alternative to RSA in public key cryptography. One advantage of ECC over RSA is key size versus strength. For example, a security strength of 80 bits can be achieved through an ECC key size of 160 bits, whereas RSA requires a key size of 1024. With a 112-bit strength, the ECC key size is 224 bits and the RSA key size is 2048 bits.

The most popular signature scheme that uses elliptic curves is called the Elliptic Curve Digital Signature Algorithm (ECDSA). The most popular key agreement scheme is called Elliptic Curve Diffie-Hellman (ECDH). An ECDH exchange is a variant of the Diffie-Hellman (DH) protocol and is an integral part of the Suite B cryptography standards proposed by the National Security Agency (NSA) for protecting both classified and unclassified information.

About Suite B

The Advanced Encryption Standard (AES) is a specification for the encryption of electronic data established by the U.S. National Institute of Standards and Technology (NIST) in 2001. Because a single encryption algorithm cannot satisfy all of the needs of the national security community, NSA created a larger set of cryptographic algorithms, called Suite B, which can be used along with AES in systems used by national security users. In addition to AES, Suite B includes cryptographic algorithms for hashing, digital signatures, and key exchanges.

Per RFC 6460, to be Suite B TLS 1.2 compliant the server and client should negotiate with the following ciphers:

- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384

RFC 6460 also lists a transitional Suite B profile for TLS 1.0 and TLS 1.1. Clients and servers that do not yet support Suite B TLS 1.2 should negotiate with the following ciphers:

- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA

There is no special configuration to ensure that Ivanti Connect Secure and Ivanti Policy Secure negotiates Suite B ciphers. However, the following general steps should be performed to enable Suite B compliance:

- An ECC certificate signed by an ECC Root CA is associated with a network port.
 - A P-256 CSR is signed by either a P-256 or P-384 Root CA.
 - A P-384 CSR is signed by a P-384 Root CA.
- Manually enable only AES128 and/or AES256 custom ciphers.

Using ECC Certificates

ECC certificates are currently supported only on the PSA hardware, PSA Series Ivanti Gateways, and virtual appliance platforms. As with RSA certificates, ECC certificates are associated with a network port. You can create multiple virtual ports on the server with each port supporting a specific certificate. For example, external virtual port 1 can use a 1024-bit RSA while external virtual port 2 uses ECC P-256 and external virtual port 3 uses ECC P-384. Only clients that support ECC cipher suites can connect to the web server on that network port.

When an Elliptic Curve Cryptography (ECC) certificate is associated with a network port, only clients that support ECC cipher suites can connect to the Web server on that network port.

Except for the key and certificate generation process, the use of ECC certificates is basically the same as using RSA certificates.

Example: Assigning an ECC P-256 Certificate to an External Virtual Port and Giving Preference to Suite B Ciphers

This example outlines the general steps for creating an external port and assigning an ECC P-256 certificate. The steps are generally the same as assigning an RSA certificate to a port.

- [Configuring the External Port](#)
- [\(optional\) Configuring the Virtual Ports](#)
- [Creating the Certificate Signing Request for a New Certificate](#)
- [Importing the Signed Certificate Created from a CSR](#)
- [Presenting the Certificate on the Network](#)
- [Setting the Security Options](#)

Configuring the External Port

The external port handles all requests from users signed into the server from outside the customer LAN, for example, from the Internet.

To configure the external port:

1. In the admin GUI, choose **System > Network > External Port > Settings**.
2. Modify the settings as needed. In this example, only IPv4 is enabled. See the following figure.

Configuring the External Port for IPv4

The screenshot displays the 'External Port' configuration page. At the top, there are tabs for 'Overview', 'Internal Port', 'External Port' (selected), 'Management Port', 'VLANs', 'Routes', 'Hosts', and 'VPN Tunneling'. Below these are sub-tabs for 'Settings', 'Virtual Ports', 'ARP Cache', and 'ND Cache'. The main content area contains the following sections:

- Use Port:** Radio buttons for 'Enabled' (selected) and 'Disabled'.
- IPv4 Settings:** Radio buttons for 'Enable IPv4' (selected) and 'Disable IPv4'. A note states: 'Note: Changing above setting might restart some services. This restart might drop all the connections to the Pulse Connect Secure.'
 - *IP Address: 192.168.5.4
 - *Netmask: 255.255.255.240 (with a link to 'cluster wide settings')
 - *Default Gateway: 192.168.5.1 (with a link to 'cluster wide settings')
 - Note: 'If you need to specify static routes, you can do so on the Static Routes page.'
- IPv6 Settings:** Radio buttons for 'Enable IPv6' and 'Disable IPv6' (selected). A note states: 'Note: Changing above setting might restart some services. This restart might drop all the connections to the Pulse Connect Secure.'
 - Link Local Address: [empty field]
 - *IPv6 Address: [empty field]
 - *Prefix Length: 64 (1 to 128) (with a link to 'cluster wide settings')
 - *Default Gateway: [empty field] (with a link to 'cluster wide settings')
- Advanced Port:**
 - MAC Address: 00:18:7D:1E:E8:FC
 - Link Speed: Auto (dropdown menu)
 - *ARP Ping Timeout: 3 seconds (3 to 300 seconds)
 - *MTU: 1500 bytes (Maximum Transmission Unit. If IPv6 is enabled (1280 to 1500), else (576 to 1500)).

A 'Save Changes' button is located at the bottom left of the form.

3. Click **Save Changes**.

(optional) Configuring the Virtual Ports

A virtual port is an IP alias bound to a physical port. It shares all of the network settings, except IP address, with the associated physical port. You can use virtual ports for different purposes, depending on the physical port or the VLAN on which you base the virtual port. In this example, we are configuring the virtual port on the external port to support external sign-ins. This is an optional step that shows one way of allowing multiple certificates on the device.

To configure the external virtual port:

1. In the admin GUI, choose **System > Network > External Port > Virtual Ports**.
2. Click **New Port**.

In this example, the port is named p_ecdsa256 and accepts only IPv4 addresses. See the following figure [Creating the Virtual Port on the External Port](#).

Creating the Virtual Port on the External Port

The screenshot shows the 'External Port' configuration page in the admin GUI. The 'Virtual Ports' tab is selected. The breadcrumb path is 'Network Settings > External Port > Virtual Ports > Virtual Port'. The form fields are as follows:

- Name:** p_ecdsa256 (with a tooltip: 'Name of the virtual port. Only alphanumeric characters, "-", or "_" are allowed.')
- Physical Port:** External Port (with a tooltip: 'The physical port determines all characteristics of this virtual port other than IP address')
- IPv4 Address:** 10.64.80.11
- IPv6 Address:** (empty field)

At the bottom, there are 'Save Changes' and 'Cancel' buttons. A note at the bottom left states '*indicates required field'.

3. Click **Save Changes**.

Creating the Certificate Signing Request for a New Certificate

A certificate signing request (CSR) is a message sent from an applicant to a certificate authority (CA) to apply for a digital identity certificate. You create a CSR through the admin console. When you create a CSR, a private key is created locally that corresponds to the CSR. If you delete the CSR at any point, the private key is deleted too, prohibiting you from installing a signed certificate generated from the CSR.

In this example, a CSR for an ECC P-256 certificate is requested.

To create a CSR for a new certificate:

1. In the admin console, choose **System > Configuration > Certificates > Device Certificates**.

2. Click **New CSR**.
3. Enter the required requestor information. In this example, the common name is ecc-p256.pulsesecure.net.
4. Click **ECC** and select **P-256** from the ECC Curve menu. See the following figure.

Creating an ECC P-256 Certificate Signing Request

Configuration > Certificates > New Certificate Signing Request


New Certificate Signing Request

Use this page to create a new Certificate Signing Request (CSR) to send to your Certificate Authority of choice.

Common Name: <small>(e.g., secure.company.com)</small>	<input type="text" value="ec-p256.psecure.net"/>
Organization Name: <small>(e.g., Company Inc.)</small>	<input type="text" value="psecure.net"/>
Org. Unit Name: <small>(e.g., IT Group)</small>	<input type="text" value="iX Group"/>
Locality: <small>(e.g., SomeCity)</small>	<input type="text" value="Sunnyvale"/>
State (fully spelled out): <small>(e.g., California)</small>	<input type="text" value="CA"/>
Country (2 letter code): <small>(i.e., US)</small>	<input type="text" value="US"/>
Email Address:	<input type="text" value="myemail@sgbix.com"/>
Key Type:	<input type="radio"/> RSA <input checked="" type="radio"/> ECC
ECC Curve:	<input type="text" value="P-256"/>

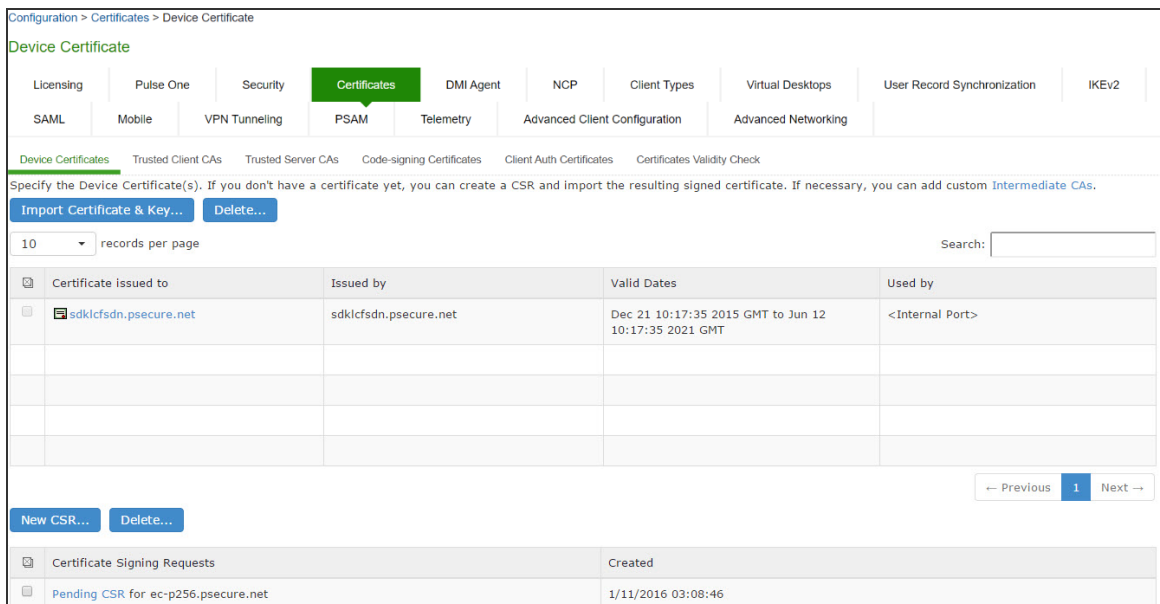
5. Click **Create CSR**. A CSR is successfully created, as shown in the following figure.

CSR Successfully Created



6. The CSR is encoded and can be copied or saved to a file. The ECC certificate should be signed by an ECC CA for Suite B compliance. Follow your CA's process for sending a CSR.
7. Click the **Back to Device Certificates** link. Until you import the signed certificate from your CA, your CSR is listed as **Pending**. See Figure *Pending CSR*.

Pending CSR



Importing the Signed Certificate Created from a CSR

Once your CA has sent your signed certificate, you must import that into the pending CSR.

To import a signed device certificate created from a CSR:

1. In the admin console, choose **System > Configuration > Certificates > Device Certificates**.
2. Under Certificate Signing Requests, click the **Pending CSR** link that corresponds to the signed certificate. See the above Pending CSR figure.
3. Under Import signed certificate, browse to the certificate file you received from the CA and then click **Import**. See the above CSR Successfully Created figure.

Presenting the Certificate on the Network

You can present a certificate many ways, depending on your configuration. For example, you can present the certificate to one or more virtual ports or on an internal or external port. In this example, the ECC P-256 certificate is presented on the external virtual port p1.

To present a certificate on an external virtual port:

1. In the admin console, select **System > Configuration > Certificates > Device Certificates**.
2. Click the certificate name you want to assign to a port. In this example, click ecc-p256.pulsesecure.net.
3. Under External Ports, select **p_ecdsa256** and click **Add**. See the following figure.

Associating the ECC P-256 with the External Virtual Port p_ecdsa256

Certificates > Certificate Details

Certificate Details

▼ **Certificate**

Issued To: ▶ sdklcfdsn.psecure.net
Issued By: ▶ ??
Valid: Dec 21 10:17:35 2015 GMT to Jun 12 10:17:35 2021 GMT
Details: ▶ Other Certificate Details
[Download](#)

▼ **Present certificate on these ports**

Select the internal and external virtual ports that will present this certificate

Internal Virtual Ports:	Selected Virtual Ports:
<div style="border: 1px solid #ccc; height: 100px; width: 100%;"></div>	<div style="border: 1px solid #ccc; height: 100px; width: 100%;"><p><Internal Port></p></div>
Add -> Remove	
External Virtual Ports:	Selected Virtual Ports:
<div style="border: 1px solid #ccc; height: 100px; width: 100%;"><p><External Port></p></div>	<div style="border: 1px solid #ccc; height: 100px; width: 100%;"></div>
Add -> Remove	
Vlan Ports:	Selected Vlan Ports:
<div style="border: 1px solid #ccc; height: 100px; width: 100%;"></div>	<div style="border: 1px solid #ccc; height: 100px; width: 100%;"></div>
Add -> Remove	

Management Port

[Save Changes](#) [Renew Certificate...](#)

4. Click Save Changes.

Setting the Security Options

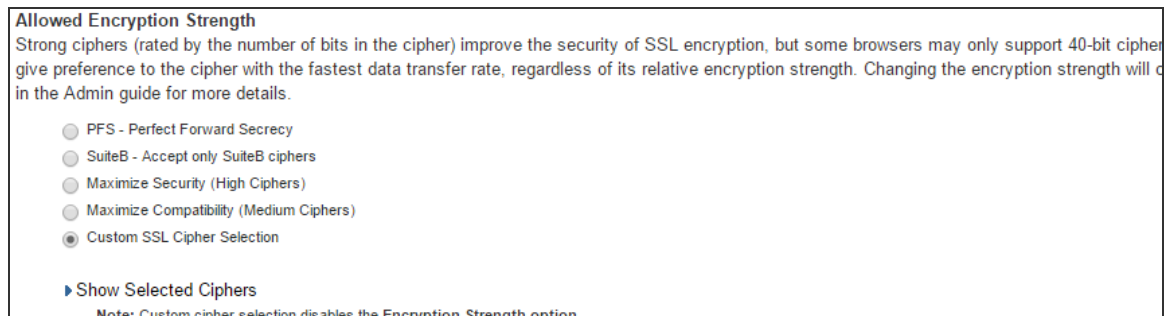
To specify the cipher suites for the incoming connection to the Web server, use the SSL Options page and select the Custom SSL Cipher Selection option. This step is required in our example to give Suite B cipher suites preference. If you do not want to give Suite B cipher suites preference, you do not have to perform this step.

Only when FIPS mode is turned on, the FIPS compliant ciphers are available to be chosen from the Supported Ciphers panel. FIPS mode is editable only on the inbound option page.

To set the security options with Inbound SSL Options:


1. In the admin console, select **System > Configuration > Security > Inbound SSL Options**.
2. Under Allowed Encryption Strength choose **Custom SSL Cipher Selection**. See the following Setting Custom SSL Cipher Selections figure.

Setting Custom SSL Cipher Selections




3. The two panels of Supported Ciphers and Selected Ciphers are displayed. Supported ciphers has the entire list of ciphers supported for the selected SSL or TLS version. Selected ciphers list the currently selected ciphers list. The below figure shows the two panels (Supported Ciphers and Selected Ciphers). Note that the Selected Ciphers and Supported Ciphers List will also be displayed for all Preset like PFS or SuiteB or Medium or High.

Supported Ciphers and Selected Ciphers Panels



4. To add a cipher to be used in order to secure a connection, click on the cipher string on the left panel and then click on the **Add>** or double click on the cipher name in the left panel. See the Setting Custom SSL Cipher Selections figure underneath.
5. To remove the cipher, click on the cipher name on the right panel and then click on the **<Remove** button or double click on the cipher name on the right side. See the Setting Custom SSL Cipher Selections figure underneath.
6. The selected ciphers on the right are listed in order of their priority from top to bottom. To change the priority of the ciphers, click on the cipher name and then click on **Move Up** to increase priority or the **Move Down** button to decrease the priority. See the Setting Custom SSL Cipher Selections figure underneath.

Setting Custom SSL Cipher Selections



7. If you are using client certificate authentication (Ivanti Connect Secure only):
 1. Select **Enable client certificate on the external port** under ActiveSync Client Certificate Configuration. See the ActiveSync Client Certificate Configuration figure underneath.
 2. Move **p_ecdsa256** to the Selected Virtual Ports column.

ActiveSync Client Certificate Configuration

Common options

SSL Handshake Timeout option
By default, the SSL handshake has a timeout of 60 seconds. Use the text box below to set a different value.

Timeout: seconds 10-600 seconds

ActiveSync Client Certificate Configuration
Enforce client certificate requirement on ports used for access. Client certificate can be enabled on the external port and/or the virtual ports.

Enable client certificate on the **external** port

8. Click **Save Changes**.

A list of the custom ciphers to be used on the device's port is displayed in the order the web server will select them. Note that Suite B ciphers are listed on top. See Confirming Custom Ciphers figure underneath. End users who now log in to external virtual port p_ecdsa256 must have at least one of the listed ciphers installed on their browser or else they cannot log in to the server.

Confirming Custom Ciphers

```

Configuration > Security > SSL Options
SSL Options

? Confirm Change to Custom Ciphers
Are you sure you want to use custom ciphers selection? If your browser does not support at least one of the ciphers listed below, you will not be able to continue.

TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA
TLS_RSA_WITH_AES_256_CBC_SHA256
TLS_RSA_WITH_AES_256_CBC_SHA
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA
TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA
TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA
TLS_RSA_WITH_3DES_EDE_CBC_SHA
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA
TLS_RSA_WITH_AES_128_CBC_SHA256
TLS_RSA_WITH_AES_128_CBC_SHA
TLS_RSA_WITH_AES_256_GCM_SHA384
TLS_RSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
TLS_DHE_RSA_WITH_AES_256_CBC_SHA
TLS_DHE_RSA_WITH_AES_128_CBC_SHA

```

9. Click **Change Allowed Encryption Strength**.

- When custom ciphers are selected, there is a possibility that some ciphers are not supported by the web browser. Also, if any of ECDH/ECDSA ciphers are selected, they require ECC certificate to be mapped to the internal/external interface. If ECC certificate is not installed, admin may not be able to log in to the box. The only way to recover from this situation is to connect to the system console and select option 8 to reset the SSL settings from the console menu. Option 8 resets the SSL settings to its default. So, the previously set SSL settings are lost. This is applicable only to Inbound SSL settings.
- Pulse Mobile client will not be able to connect to ICS device, if the ciphers selected in Inbound option are not supported by the mobile client.



Enabling Outbound SSL Options

Only for Outbound SSL Settings, we can configure Non FIPS Ciphers when FIPS is Enabled using Custom Cipher Selection Option. Now, there are options to change different SSL/TLS versions and different encryptions in the Outbound SSL Settings. [Outbound SSL Settings](#) shows the Outbound SSL Settings.

Outbound SSL Settings

Verifying the Certificate on the Client

End users can check which certificate their browser is using to connect to the server. In the following example, the end user connects to server port p3, which uses an ECC curve P-256 certificate. See the following figure.

Connecting to a Port Using an ECC Curve P-256 Certificate

Certificate issued to	Issued by	Valid Dates	Used by
ps.ivanti.com	ps.ivanti.com	Mar 8 10:33:55 2022 GMT to Aug 29 10:33:55 2027 GMT	<Internal Port>

To view the certificate from an Internet Explorer 8 browser:

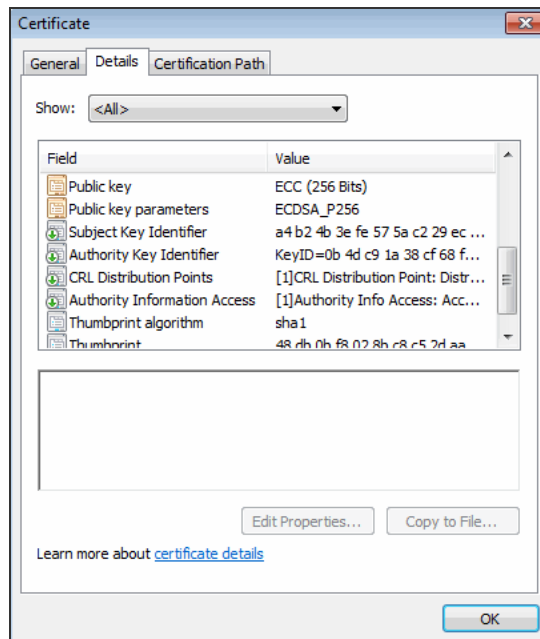
1. Open an Internet Explorer 8 browser and point to the server to which you want to connect.
2. Click the lock icon located at the end of the address bar and then click the **View Certificate** link. See the following figure.

Viewing the Connection Certificate Information



3. Click the **Details** tab and scroll down until you see the Public key field. In this example, the public key value is ECC (256 Bits) which matches the server port p3 certificate shown in the following figure.

Certificate Public Key



Using TCP Dump to View Cipher Information

You can use the TCP Dump tool to view which cipher each client uses to connect to the server. TCP Dump is a packet analyzer that intercepts (sniffs) and displays TCP/IP and other packets transmitted or received between the server and clients.



To permit debugging, it is recommended that the ECC certificate be replaced by an RSA certificate so that an RSA cipher suite gets selected and then the application data can be decoded.

To capture packet headers:

1. Select **Maintenance > Troubleshooting > Tools > TCP Dump**.
2. Select the interface, internal or external or both, you wish to sniff and then the VLAN port.
3. Click **Start Sniffing**.


The next time a user points a browser window to the server or logs in to the server, handshake information is obtained.

4. Click **Stop Sniffing** when done.

To view the packet headers:

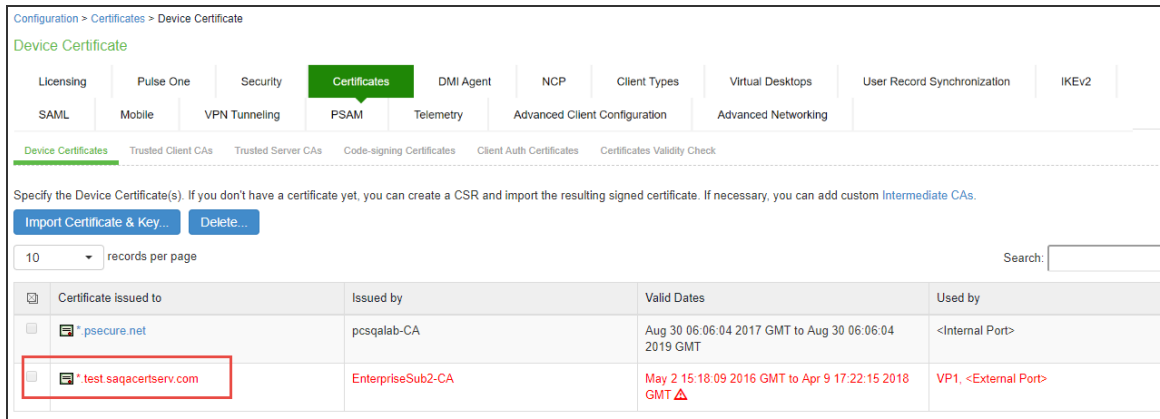
1. Select **Maintenance > Troubleshooting > Tools > TCP Dump**.
2. Under Dump file, select **SSLDump** from the file menu and the certificate to use. See the figure underneath.

Viewing the TCP Dump Output



The certificate names in the TCP Dump window are the same as the "Certificate issued to" names in the Device Certificates window. Select the certificate corresponding to the port you wish to view packet information.

Issued to Certificate on the Device Certificates Pages (**System > Configuration > Certificates > Device Certificate**):



3. Click **Get**.

Portions of a TCP dump output follow.

The client starts a handshake with the server:

```
1 1 0.0007 (0.0007) C>S Handshake
```

The client then lists its supported cipher suites:

```
cipher suites
    TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384
    TLS_ECDH_ECDSA_WITH_AES_256_SHA384
    TLS_ECDH_ECDSA_WITH_AES_256_SHA
    TLS_ECDH_ECDSA_WITH_DES_CBC3_SHA
    ...
```

The server acknowledges the handshake:

```
1 2 0.0010 (0.0003) S>C Handshake
```

The server compares the cipher suites on the client with the ones on the server and picks the cipher suite that is preferred by the server based on SSL options:

```
cipherSuite      TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384
```

Example TCP Dump Output

New TCP connection #1: 10.64.8.3(46200) <-> 10.64.90.21(443)

```
1 1 0.0007 (0.0007) C>S Handshake
    ClientHello
        Version 3.3
        cipher suites
            TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384
```

```

    TLS_ECDH_ECDSA_WITH_AES_256_SHA384
    TLS_ECDH_ECDSA_WITH_AES_256_SHA
    TLS_ECDH_ECDSA_WITH_DES_CBC3_SHA
    TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA384
    TLS_ECDH_ECDSA_WITH_AES_128_SHA256
    TLS_ECDH_ECDSA_WITH_AES_128_SHA
    TLS_ECDH_ECDSA_WITH_RC4_SHA
    Unknown value 0xc001
    TLS_EMPTY_RENEGOTIATION_INFO_SCSV
    compression methods
        NULL
    ClientHello Extensions [113]=
        00 6f 00 0b 00 04 03 00 01 02 00 0a 00 34 00 32
        00 0e 00 0d 00 19 00 0b 00 0c 00 18 00 09 00 0a
        00 16 00 17 00 08 00 06 00 07 00 14 00 15 00 04
        00 05 00 12 00 13 00 01 00 02 00 03 00 0f 00 10
        00 11 00 23 00 00 00 0d 00 22 00 20 06 01 06 02
        06 03 05 01 05 02 05 03 04 01 04 02 04 03 03 01
        03 02 03 03 02 01 02 02 02 03 01 01 00 0f 00 01
        01
1 2 0.0010 (0.0003) S>C Handshake
    ServerHello
        Version 3.3
        session_id[32]=
            a3 07 40 6e 73 12 c2 4d f3 7d b9 77 f8 97 e1 94
            fc 1b 51 6a 66 3c 99 d6 c7 7d 0e fa 29 2e d0 c4
        cipherSuite          TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384
        compressionMethod    NULL
        ServerHello Extensions [20]=
            00 12 ff 01 00 01 00 00 0b 00 04 03 00 01 02 00
            0f 00 01 01
1 3 0.0010 (0.0000) S>C Handshake
    Certificate
1 4 0.0010 (0.0000) S>C Handshake
    ServerHelloDone
1 5 0.1413 (0.1403) C>S Handshake
    ClientKeyExchange
1 6 0.1413 (0.0000) C>S ChangeCipherSpec
1 7 0.1413 (0.0000) C>S Handshake
1 8 0.1464 (0.0051) S>C ChangeCipherSpec
1 9 0.1464 (0.0000) S>C Handshake
1 10 9.2389 (9.0924) C>S application_data

```

```
1 11 9.5828 (0.3438) C>S application_data
1 12 9.5833 (0.0004) S>C application_data
1   9.5833 (0.0000) S>C TCP FIN
1 13 9.5999 (0.0166) C>S Alert
1   9.5999 (0.0000) C>S TCP FIN
```

Configuration File Administration

Configuration File Administration Overview

The system supports multiple administrator utilities related to configuration file management. The following table describes the purpose of the different utilities.

The following table lists the Utilities for Configuration File Administration:

Utility	Recommended Usage
Archiving	Schedule periodic backups to a remote backup server. You should schedule archiving for both the system configuration binary file (system.cfg) and the user configuration binary file (user.cfg). If necessary, you can import an archived configuration using the configuration binary file import/export feature.
Local backup and restore	Create backups on the local system as a precaution when making significant configuration changes. With this utility, you can quickly restore to a previous configuration.
Binary configuration file import/export	<p>Export binary configuration files to a local host (an alternative to the remote archiving server and archiving process that runs as a scheduled job). You might do this if you do not use or do not have access to an archiving server, or if you want to make use of a configuration that has not yet been archived. You can export the binary system configuration file (system.cfg) and the binary user configuration file (user.cfg).</p> <p>You can use the binary file import/export feature to clone a configuration that you want to deploy more broadly, such as deploying a backup device or to a group of devices. You can use "selective import" options to exclude unique network identifiers (such as IP address) that would cause problems if the configuration were to be wholly imported and activated.</p>

Utility	Recommended Usage
XML configuration file import/export	<p>Import or export the configuration for only the features and settings you select. This enables you to take a more granular approach to mass configuration management than the binary file import/export feature. For example, you might want to populate an authentication server configuration across a large number of nodes. You can export just that configuration element, and when you import it in the other nodes, you do not overwrite the large number of configuration elements that you would if you had imported the user.cfg file.</p> <p>You might also find the XML file import/export feature useful when managing a single node. For example, you might want to add many new users to the local authentication server, which can be faster editing the XML than using the user interface. Or you might want to make global changes to the configuration object naming conventions or descriptions as part of a "housekeeping" initiative. This, too, might be accomplished faster editing the XML than clicking through the user interface.</p>
Push configuration	<p>Push a partial configuration from the running configuration on the source system to the running configuration on one or more target systems. This is the best option to instill common configuration elements if the devices are already deployed and currently online.</p>

Configuring Archiving for System Logs, Configuration Files, and Snapshots

You can schedule periodic archiving for system logs, system configuration files, and system snapshots. Periodic archiving occurs only at the scheduled time. "Unscheduled" archiving does not occur automatically. For example, if a log file exceeds the maximum file size, the archiving process does not automatically back up the file prior to the scheduled time to prevent data loss.

If the archive process fails, it makes two more attempts at an interval of 30 seconds. If the archiving still fails, it retries at an interval of one hour till the archiving process is successful.

We recommend that you schedule an archive operation during hours when traffic is light to minimize its impact on users. The automatic archiving process compresses files and, if the system is busy, can degrade performance for users. Also, a cluster node might appear unresponsive if the system is busy with traffic and performing archiving simultaneously.



If you schedule an archive operation to occur during the hour that your system switches to daylight savings time (DST), the operation might not occur as scheduled. For example, if your system is set to change to DST at 1:00 AM, and you scheduled an archive job to occur at any time between 1:01 AM and 1:59 AM, then the operation does not take place because at 1:00 AM the system clock is moved forward to 2:00 AM, and the system never reaches the archive time for that date.

To configure log archiving:

1. Select **Maintenance > Archiving > Archiving Servers** to display the configuration page.

[Archiving Configuration Page - Ivanti Connect Secure](#) shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration as described in [Archiving Configuration Page - Ivanti Connect Secure](#).
3. Save the configuration.

Archiving Configuration Page - Ivanti Connect Secure

Archiving > Archiving Servers

Archiving Servers

Archiving Servers Local Backups Pulse Collaboration

You can schedule automatic archiving of log data, system configuration, and user accounts. To do so, specify accessible location for the data, an account to use, and the specific schedule for each type of archive data.

Archive Settings

Method: SCP FTP AWS S3 Azure Storage

*S3 Bucket Name: AWS S3 bucket name

*Region: AWS S3 bucket location

*AWS Access Key: AWS account access key

*AWS Secret Key: AWS account secret key

Dest Path Prefix: Path to copy files under S3 bucket, eg: folder1/folder2

* indicates required field

Archive Schedule

Select one or more components to schedule an archive.

Archive events log

Archive user access log

Archive admin access log

Archive Sensors log

Archive client-side log uploads

Archive system configuration

Archive user accounts

Archive Administrative Network Configuration

Archive XML configuration

Archive Debug Log

Archive Periodic SnapShots

The following table lists the Archiving Configuration Guidelines:

Settings	Guidelines
Archive Settings	
Archive Server	Specify the fully qualified domain name or IP address of the server to which to send the archive files.
Destination Directory	Specify the destination directory. Follow these recommendations: For UNIX systems, you can specify an absolute or relative path. We recommend you specify a full path. For Windows systems, specify a path that is relative to the ftproot directory. We recommend you specify a full path. Do not include a drive specification for the destination directory, such as: ivanti/log.
Username	Specify a username that has privileges to log into the server and write to the destination directory.
Password	Specify the corresponding password.
Method	Select SCP, FTP, AWS S3 or Azure Storage . SCP is the default method. SCP is a file transfer utility similar to FTP. SCP encrypts all data during transfer. When the data reaches its destination, it is rendered in its original format. SCP is included in most SSH distributions and is available on all major operating system platforms. AWS S3: Push backup configurations and archived logs to Amazon AWS S3 bucket. For more details, refer to Ivanti Connect Secure Virtual Appliance on Amazon AWS Cloud Deployment Guide. Azure Storage: Push backup configurations and archived logs to Microsoft Azure storage. For more details, refer to Ivanti Connect Secure Virtual Appliance on Microsoft Azure Cloud Deployment Guide.
Archive Schedule	

Settings	Guidelines
Archive events log	<p>Schedule archiving for the Events log. The archive file has the following format:</p> <p>PulseSecureEventsLog-[clustername standalone]-[nodename hostname]-[date]-[time]</p> <p>For example, an archive file for a cluster named Gen has a filename similar to the following: PulseSecureEventsLog-Gen-node1-Root-20090109-1545.gz.</p> <p>The archiving schedule configuration includes the following options:</p> <p>Use this filter-Select a log format filter.</p> <p>Day of week-Select the days of the week on which to run the archiving job. Every hour or a Specified Time. The Every hour option runs a job every hour on the hour for the selected days. The specified time option runs a job once on the selected days.</p> <p>Clear log after archiving. Select this option to clear the local log file after the archiving job is successfully completed. If an archive job fails, the log files are not deleted.</p> <p>Password-(Optional) Specify a password to secure and encrypt system configuration or user account archives.</p>
Archive user access log	<p>Schedule archiving for the User Access log. The archive file has the following format:</p> <p>PulseSecureAccessLog-[clustername standalone]-[nodename hostname]-[date]-[time]</p> <p>The archiving schedule configuration includes the same options as those described for the Events log.</p>
Archive admin access log	<p>Schedule archiving for the Admin Access log. The archive file has the following format:</p> <p>PulseSecureAdminLog-[clustername standalone]-[nodename hostname]-[date]-[time]</p> <p>The archiving schedule configuration includes the same options as those described for the Events log.</p>
Archive client-side log uploads	<p>Schedule archiving for client-side log uploads. This option is available only on Ivanti Connect Secure.</p> <p>The archiving schedule configuration includes the same options as those described for the Events log, except for log filter format, which is not applicable to the client-side logs.</p>

Settings	Guidelines
Archive system configuration	<p>Schedule archiving for the system configuration binary file (system.cfg). The archive file has the following format:</p> <p>PulseSecureConf-[clustername standalone]- [nodename hostname]- [date]- [time]</p> <p>The archiving schedule configuration includes the same day, time, and password-protection options as those described for the Events log.</p>
Archive user accounts	<p>Schedule archiving for user account configuration binary file (user.cfg). The archive file has the following format:</p> <p>PulseSecureUserAccounts-[clustername standalone]- [nodename hostname]- [date]- [time]</p> <p>The archiving schedule configuration includes the same day, time, and password-protection options as those described for the Events log.</p>
Archive XML configuration	<p>Schedule archiving for the XML configuration files.</p> <p>The archiving schedule configuration includes the same day and time options as those described for the Events log.</p> <p>Administrator can select the Exclude large-data packages in the archived configuration to exclude ESAP and Pulse Client packages from being archived.</p>
Archive debug log	<p>Enable archiving for collected debug logs.</p> <p>You cannot specify a day and time for archiving debug logs. If you select this option, debug logs are archived periodically and cleared if the Clear log after archiving option is selected.</p>
Archive periodic snapshots	<p>Enable archiving for snapshots.</p> <p>You cannot specify a day and time for archiving periodic snapshots. If you select this option, snapshots are archived periodically.</p>

Using the Configuration Backup and Restore Feature

You can save up to five system configuration backups and five user account backups on the local server. If you exceed this limit, the system overwrites the oldest backup with the new backup. If you do not want to overwrite the oldest backup, select and delete another backup instead, before you save the most current one.


To manage configuration file backups:

1. Select **Maintenance > Archiving > Local Backups** to display the configuration page.

The figure underneath shows the archiving configuration page for Ivanti Connect Secure.

- 2. Use the controls to backup or restore the configuration as described in the following table.
- 3. Save the configuration.

Local Backups Management Page - Ivanti Connect Secure



The following table lists the Local Backups Management Guidelines:

Controls	Guidelines
System Configuration	
Save Configuration	Create a backup of the running system configuration.
Delete	Select a row in the table and click Delete to delete the backup.
Restore	Select a row in the table and components in the "Include when restoring" column and click Restore to replace the running configuration with the archived configuration.

Controls	Guidelines
User Configuration	
Save User Accounts	Create a backup of the running user configuration.
Delete	Select a row in the table and click Delete to delete the backup.
Restore	Select a row in the table and click Restore to replace the running configuration with the archived configuration.

Using the Import/Export Feature for Binary System Configuration Files

This topic describes the import/export feature for binary system configuration files.

Binary System Configuration File Overview

The Ivanti access management framework enables you to import and export the system and network settings using binary system configuration files. When importing a system configuration file, you can exclude the device certificate and the server's IP address or network settings from the imported information. For example, to set up multiple Ivanti Connect Secure systems behind a load balancer, import everything except for the IP address. To set up the system as a backup server, import everything except for the digital certificate and the network settings.

The binary system configuration file includes the following settings:

- Network settings
- Certificates. The system imports only device certificates, not the chains that correspond to the device certificates or trusted client CAs.
- Cluster configuration
- Licenses. When you import a configuration file that contains licenses, the system gives precedence to any existing licenses. Licenses are imported only if no licenses are currently installed.
- SNMP settings
- Client-side logs. To import or export client-side logs, import or export both the system and user configuration files.

- Web proxy servers. Ivanti Connect Secure only. To export all web proxy related information, both the system and user configuration files are needed.
- Web caching options. Ivanti Connect Secure only.
- Rewriter filters. Ivanti Connect Secure only.

Exporting a Binary System Configuration File

To export a binary system configuration file:

1. Select **Maintenance > Import/Export > Import/Export Configuration** to display the configuration page.

The figure underneath shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration and export operation as described in [Table](#).

Export Binary System Configuration File Configuration Page - Ivanti Connect Secure

Import/Export > System Configuration

System Configuration

Configuration | User Accounts | XML Import/Export

Export

To export system settings to a configuration file, click Save Config As. You can optionally password-protect this file:

Password for configuration file:

Confirm Password:

Save Config As...

Import

To import system settings from a configuration file, select the configuration file and which settings to bring in, and click Import Configuration Options:

Import Device Certificate(s)
Note: Checking this will overwrite the existing Device Certificate(s).

Other Import Options:

- Import everything (except Device Certificate(s))
- Import everything but the IP address
 Preserves the IP address, netmask, default gateway, VIPs, ARPs and routes of the network interfaces on this device.
Note: Use this option only if the exported configuration file is from a standalone node.
- Import everything except network settings, cluster settings and licenses
 Leaves everything in Network Settings, Clustering Properties, Licensing sections and Onboarding Profile UUID unchanged.
Note: Always use this option if configuration file was exported from a node that is part of a cluster.
- Import only Device Certificate(s)
 Imports the Device Certificate(s) only.
Note: You must check the Import Device Certificate(s) checkbox above.

Config File: **Browse** No file chosen

Password: Use this if the configuration file was password-protected

Note that importing configuration with a different SSL acceleration setting will reboot the IVE.

Import Config

The following table lists the Export Binary System Configuration File Configuration and Action Guidelines:

Settings	Guidelines
Password for configuration file	Specify a password to encrypt and secure the configuration file.
Confirm password	Specify the password.
Save Config As	Display a dialog box to save the file to your local host.

Importing a Binary System Configuration File

To import a binary system configuration file:

1. Select **Maintenance > Import/Export > Import/Export Configuration** to display the configuration page.

The following figure shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration and import operation as described in the following table.

Import Binary System Configuration File Configuration Page:

The screenshot shows the 'System Configuration' page with three tabs: 'Configuration' (selected), 'User Accounts', and 'XML Import/Export'. Under the 'Export' section, there are fields for 'Password for configuration file:' and 'Confirm Password:', followed by a 'Save Config As...' button. The 'Import' section includes a checkbox for 'Import Device Certificate(s)?' with a note that it will overwrite existing certificates. Below this are radio button options for 'Other Import Options': 'Import everything (except Device Certificate(s))', 'Import everything but the IP address' (with a note about preserving IP settings), 'Import everything except network settings, cluster settings and licenses' (with a note about leaving network settings unchanged), and 'Import only Device Certificate(s)' (with a note that it imports only certificates). At the bottom, there is a 'Config File:' field with a 'Browse' button and 'No file chosen' text, and a 'Password:' field with a note to use it if the file is password-protected. A final note states that importing with a different SSL acceleration setting will reboot the IVE, followed by an 'Import Config' button.

The following table lists the Import Binary System Configuration File Configuration and Action Guidelines:

Settings	Guidelines
Options	

Settings	Guidelines
Import Device Certificate(s)?	<p>Overwrite the existing device certificate(s) with the ones in the imported configuration file.</p> <p>When importing a device certificate in to a FIPS device, note that you must choose a certificate that uses a FIPS-compliant private key. To ensure FIPS-compliance, select a certificate and corresponding security world private keys were generated on a FIPS device.</p>
Other Import Options	
Import everything (except Device Certificate(s))	Import all settings except the device certificate.
Import everything but the IP address	<p>Do not overwrite the existing configuration for network interface IP addresses, netmask, default gateway, virtual interfaces, ARP tables, and route tables. Use this option only if the exported configuration file is from a standalone node.</p> <p><i>To set up multiple nodes in a cluster behind a load balancer, import everything except the IP address.</i></p>
Import everything except network settings and licenses	<p>Do not allow the imported configuration to change the existing configuration for settings found in the Network Settings and Licensing sections. With this option, network configurations, licenses, cluster configurations, certificates, defined SNMP settings and syslog configurations are not imported. Always use this option if configuration file was exported from a node that is part of a cluster.</p> <p><i>To set up a backup node, import everything except network settings and digital certificates.</i></p>
Import only Device Certificate(s)	Import the device certificate(s) only. You must also select the Import Device Certificate(s) check box.
Config file	Use the browse button to locate and select the file from your local host.
Password	Specify the password (if applicable).
Import Config	Import the configuration file.

Using the Import/Export Feature for Binary User Configuration Files

This topic describes the import/export feature for user configuration binary files.

Binary User Configuration File Overview

In general, if a menu item falls under the Authentication, Administration, or Users menu, the item is included in the user configuration file (user.cfg). In particular, the user configuration file includes the following settings:

- Sign-in settings (includes sign-in policies, sign-in pages, all authentication servers, authentication protocol sets, and Ivanti settings)
- Authentication realms (including admin realms, user realms, and MAC authentication realms)
- Roles
- Resource profiles. Ivanti Connect Secure only.
- Resource policies
- User accounts
- Client-side logs. To export or import client-side logs, export or import both the system and user configuration files.

Exporting a Binary User Configuration File


To export a binary user configuration file:

1. Select **Maintenance > Import/Export > Import/Export Users** to display the configuration page.

The figure underneath shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration and export operation as described in [Table](#).

Binary Export User Configuration File Configuration Page - Ivanti Connect Secure



The following table lists the Binary Export User Configuration File Configuration and Action Guidelines:

Settings	Guidelines
Password for configuration file	(Optional) Specify a password to encrypt and secure the configuration file.
Confirm password	Specify the password.
Save Config As	Display a dialog box to save the file to your local host.

Importing a Binary User Configuration File


To import a binary user configuration file:

1. Select **Maintenance > Import/Export > Import/Export Users** to display the configuration page.

The following figure shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration and import operation as described in the following table.

Import User Configuration Binary File Configuration Page:



The following table lists the Import Binary User Configuration File Configuration and Action Guidelines:

Settings	Guidelines
Browse	Locate and select the file from your local host.
Password	Specify the password (if applicable).
Import Config	Import the configuration file.

Using the Import/Export Feature for XML Configuration Files

This topic describes the import/export feature for XML configuration files.

XML Configuration File Overview

The system maintains its configuration in a structured XML file. This enables the system to support an alternative to the complete configurations that are exported and imported with the configuration binary files. You can use the export/import configuration XML pages to export and import selected configuration elements.

You might find the feature useful when performing the following tasks:

- Adding to the configurations of peer nodes, for example, adding a large number of users.
- Modifying multiple instances of a single setting, for example, an authentication server name.
- Deleting settings, for example, deleting authentication servers that are no longer used.
- Creating a configuration template to use for setting up new nodes.
- Tracking configuration changes by comparing differences on periodic exports.

Guidelines and Limitations

The following table summarizes the guidelines and limitations for using the XML import/export feature.

The following table lists the XML Import/Export Guidelines and Limitations:

Category	Guidelines and Limitations
General	<p>The following guidelines and limitations apply:</p> <p>You can import and export configuration files only between systems running the same software version.</p> <p>If XML configuration to be imported contains one or more Pulse Client packages, we recommend to split the configuration to import only Pulse Client packages first considering one Pulse Client package per import and then import the remaining configurations.</p> <p>You might find it useful to use a text editor to modify configuration elements that ought to be distinguished, such as configuration object names and descriptions. Never modify the names of the NIC identifiers. The system relies on knowing that each appliance has two interface cards, known as NIC0 and NIC1.</p> <p>Immediately after importing an Active Directory authentication server configuration, you must edit the configuration to change the Computer Object name. Unexpected problems might arise if two systems join an Active Directory domain using the same Computer Object name.</p>

Category	Guidelines and Limitations
Licenses	<p>The following rules apply to exported and imported licenses:</p> <p>You cannot edit the license data that is exported. It is encrypted.</p> <p>An XML import of licenses is valid only if the system does not currently have a license installed. If a license is installed already, any imported licenses are dropped. If you still intend to import a license, you must perform a factory reset before you perform the import operation.</p> <p>If you import a license after deleting a temporary license, the imported license is dropped because you might still be able to reactivate the deleted license. The import operation preserves any licensing data.</p>
Clusters	<p>The following guidelines apply to importing a configuration file for nodes that belong to a cluster:</p> <p>When you perform an import operation on a cluster, all of the cluster nodes must be enabled and running. If you attempt to import a configuration into a cluster in which a node is not running, the import operation might hang or your import results might be unpredictable.</p> <p>The XML configuration that you import must contain the same set of nodes as the original cluster. The signature used to synchronize the cluster when the nodes are reenabled is derived from the IP addresses of the cluster nodes. Therefore, the remaining nodes cannot rejoin the cluster if the imported configuration yields a different signature.</p> <p>When import occurs, the imported configuration file overwrites the node-specific cluster configuration network settings of the remaining nodes. If you change the node-specific network settings, make sure you do not make the remaining nodes unreachable.</p> <p>After you have exported the file, do not modify settings that could render the primary node unreachable, such as changes to network settings.</p> <p>After you have exported the file, do not modify the XML to change the node name, IP address, or IP netmask.</p> <p>After you have exported the file, do not modify virtual port settings or add new virtual port settings.</p>

Exporting an XML Configuration File

To export an XML configuration file:

1. Select **Maintenance > Import/Export > Export XML** to display the configuration page.

[Export XML File Configuration Page - Ivanti Connect Secure](#): shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration and export operation as described in [Table](#).

Export XML File Configuration Page - Ivanti Connect Secure:

Import/Export > Export XML

Export XML

Configuration | User Accounts | **XML Import/Export**

Export | Import

▼ **Schema Files**

Download the [Schema files](#)

▼ **Select Settings and Export**

[Expand All](#) [Select All](#) [Export...](#)

- ▶ **System Settings...** none selected
- ▶ **Sign-in Settings...** none selected
- ▶ **Endpoint Security...** none selected
- ▶ **Authentication Realms...** none selected
- ▶ **Roles...** none selected
- ▶ **Resource Profiles...** none selected
- ▼ **Resource Policies...** none selected
 - Select All Resource Policies**
 - Note that resource policies related to Resource Profiles cannot be exported.
 - Web**
[All](#) | [None](#)
 - Access Control
 - Java access control
 - Selective rewriting
 - Custom Headers
 - Remote SSO form post
 - SAML SSO
 - Web proxy
 - ActiveX parameter rewriting
 - HTTP Protocol
 - Cross Domain Access
 - SAML External Apps SSO
 - Caching
 - Java code-signing
 - Passthrough proxy
 - SSO Basic Auth/NTLM/Kerberos
 - Remote SSO headers/cookies
 - SAML access control
 - Launch JSAM
 - Compression
 - Web Encoding
 - Client Authentication
 - Options
 - Files**
[All](#) | [None](#)
 - Windows access control
 - Windows compression
 - UNIX/NFS compression
 - Windows Credentials
 - UNIX/NFS access control
 - Options
 - SAM Applications**
[All](#) | [None](#)
 - SAM access control
 - Options
 - Telnet/SSH**
[All](#) | [None](#)
 - Telnet/SSH access control
 - Options
 - VPN Tunneling**
[All](#) | [None](#)
 - Access control
 - Split tunneling networks
 - Connection profiles
 - Bandwidth Management
 - Terminal Services**
[All](#) | [None](#)
 - Terminal Services access control
 - Options
 - Email Client**
[All](#) | [None](#)
 - Email Settings
 - HTML5 Access**
[All](#) | [None](#)
 - HTML5 Access access control
 - Options
- ▶ **Pulse Secure client...** none selected
- ▶ **Enterprise Onboarding...** none selected
- ▶ **Local User Accounts...** none selected
- ▶ **Maintenance Settings...** none selected

[Export...](#)

The following table lists the Exporting an XML Configuration File settings and guidelines:

Settings	Guidelines
Schema Files	
Schema files	Download the XML schema definition (.xsd) files that describe the XML.
Select Settings and Export	
Expand All	Expand the display of all settings groups.
Select All	Select all settings for all groups.
Export	Export the selected configuration data to an XML file.
Settings	
System	Expand this group and select settings found under the System menu. Do not select the DMI Agent unless Support Agent instructs you to do so.
Sign-in	Expand this group and select settings found under the Sign-in menu.
Endpoint Security	Expand this group and select settings found under the Endpoint Security menu. ESAP packages are encrypted when exported.
Authentication Realms	Expand this group and select authentication realm settings, including user realms and MAC address authentication realms.
Roles	Expand this group and select settings found under the Roles menu.
Resource Profiles	Ivanti Connect Secure only. Expand this group and select settings found under the Resource Profiles menu.
Resource Policies	Expand this group and select settings resource policies settings.
Pulse Client	Expand this group and select settings found under the Pulse client menu.
Local User Accounts	Expand this group and select local authentication server settings.
Maintenance	Expand this group and select settings found under the Maintenance menu.
Export Settings?	
Export	Export the selected configuration data to an XML file.

Importing an XML Configuration File


To import an XML configuration file:

1. Select **Maintenance > Import/Export > Import XML** to display the configuration page.

Figure underneath shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration and import operation as described in [Table](#).

Import XML File Configuration Page:



The following table lists the Import XML File Configuration and Action Guidelines:

Settings	Guidelines
Schema Files	
Schema files	Download the XML schema definition (.xsd) files that describe the XML.
Import	
XML data file	Locate and select the XML file.

Settings	Guidelines
Import	Import the file. The Import XML Results page is displayed. This page contains information about the imported network settings, roles, resource policies, and other settings. If there are errors in the XML, the import operation stops and rolls back the configuration to the previous state. Error messages are displayed on the Import XML Results page.

Example: Using the Configuration XML File Import/Export Feature to Add Multiple Users

This example shows how to use the configuration XML file import/export feature. The example is illustrative. There are additional ways to use export files.

Assume you have just added a new device to the network, and you want to add your 2,000 users to the system. Instead of adding them one at a time in the admin console, you want to perform a mass import. You can export the user accounts, extract the relevant XML that defines users, replicate each element as needed, and then import them. In this situation, your configuration should include the option to force the users to change their passwords the first time they log in to the system.

In this procedure, you only see examples for User 1, User 2, and User 2000. All other users are included in your import file. You set the passwords to numbered instances of the word password, such as password1, password2, and so on. All users in this example are assigned to the same auth server, although you can specify any combination of auth servers that are valid on your system.

To add multiple new users:

1. Select **Maintenance > Import/Export > Export XML**.
2. Follow the instructions to export local user accounts.
3. Save the exported file as users.xml.
4. Open the **users.xml file**.
5. Copy and paste the User container element repeatedly until you have added the necessary number of users. Although the example shows only three new users, you might add hundreds of new users to the file.
6. Update the appropriate data in each User container element as shown in "Example: Updating the User container".
7. Save the **users.xml file**.

8. Select **Maintenance > Import/Export > XML Import/Export > Import**.
9. Click **Browse** to locate and select your users.xml file.
10. Click **Import**.

Example: Updating the User container

```
<configuration xmlns="http://xml.juniper.net/ive-sa/6.2R1"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <authentication>
    <auth-servers>
      <auth-server>
        <local>
          <users>
            <user>
              <username>user1</username>
              <fullname>User1</fullname>
              <password-cleartext>password1
            </password-cleartext>
              <one-time-use>>false</one-time-use>
              <enabled>>true</enabled>
              <change-password-at-signin>>true
            </change-password-at-signin>
            </user>
            <user>
              <username>user2</username>
              <fullname>User2</fullname>
              <password-cleartext>password2
            </password-cleartext>
              <one-time-use>>false</one-time-use>
              <enabled>>true</enabled>
              <change-password-at-signin>>true
            </change-password-at-signin>
            </user>
          <name>System Local</name>
        </auth-server>
      </auth-servers>
    </authentication>
  </configuration>
```

Guidelines for Modifying Configuration XML Files

This topic provides guidelines for modifying an exported configuration file.

Preparing to Modify a Configuration XML File

The following practices might be useful when you export and import XML configuration elements:

- Define your goals for a particular task:
 - What object or objects do you need to add, update, or delete?
 - Do you need to complete all modifications in one operation, or can you modify the configuration in separate operations?
 - Is your process a one-time operation, or do you need to perform the same operation multiple times?
 - Are you updating an active system or are you using one configuration as a template for configuring systems that have not yet been brought online?
- Document the intended changes to the configuration objects:
 - Make a list of objects to be added, updated, or deleted.
 - For objects to add or update, list specific attribute data.
 - List pages or tabs from the admin console that correspond to the objects and attributes you intend to change.
- Make a binary system snapshot or a binary configuration backup immediately before you perform the import.
- Make a plan to verify that the completed configuration meets your goals.
 - View the Admin Access log to make sure the export and import operations succeeded.
 - Perform a random check of the modified items. Make sure items were added, updated, or deleted as you expected.
- Make sure you are able to view configuration details in both the admin console and XML file while you work on the modifications, typically in the following sequence.
 1. Use the admin console to correlate the configuration data with the data in the XML file:
 2. Use the XML file to locate and modify the configuration data.

3. Use the admin console to verify the successful import.
 - Use an XML editor. The exported XML files have a standard structure. Once you become familiar with the structure, you can navigate the files easily. The files can become large, so you might find it more efficient to use a commercial or open source XML editor. XML editors often separate the editable data from the structural display. This separation reduces or eliminates the risk of accidentally modifying an XML element rather than its data.

Understanding the XML Export File

When you export a configuration file, the system saves the configuration as an XML file. The data in the exported file is based on the selections you make when you configure the export operation. The file contains all of the required XML processing instructions and namespace declarations, which must be included exactly as defined.

The following table provides some basic information and guidelines to help you understand the structured XML used in the export file.

The following table lists the Structured XML Files: Basic Information and Guidelines:

Topic	Guideline
XML schema definition (.xsd) file	<p>The export is based on an XML schema. The schema is a separate file that defines the metadata, and that serves as a model or a template for the exported file. Use the XML schema file to:</p> <ul style="list-style-type: none"> Identify the structure and sequence of configuration objects. Identify optional and required elements, allowable values, default values, and other attributes of the configuration objects. <p>You can download the XML schema definition (.xsd) file in either of the following ways:</p> <ul style="list-style-type: none"> From the XML Import/Export pages by clicking a link. From the URL where the files are stored on the system (you do not need to sign in). <p>To access the .xsd file, access the following URL: <a href="https://<IP-or-hostname>/dana-na/xml/config.xsd">https://<IP-or-hostname>/dana-na/xml/config.xsd</p>
Elements	<p>An element is a discrete XML unit that defines an object or part of an object. The element typically consists of a pair of tags that may or may not surround string data. Tags are surrounded by angle brackets (< >).</p>

Topic	Guideline
Namespaces	<p>Namespaces allow you to use the same words or labels in your code from different contexts or XML vocabularies. Prefixing elements with namespace qualifiers allow the XML file to include references to different objects that originate in different XML vocabularies and that share the same name. If you do not prefix elements with namespace qualifiers, the namespace is the default XML namespace, and you refer to element type names in that namespace without a prefix.</p> <p>A namespace declaration looks like the following example:</p> <pre><configuration xmlns="http://xml.juniper.net/ive-sa/6.2R1" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"></pre> <p>When you see namespace identifiers in your XML files, you do not need to be concerned about them, as long as you do not delete or modify the identifiers.</p>
Element Sequence	<p>You should avoid changing the sequence of elements in the XML file, whenever possible. Although the schema does not enforce sequence in all cases, you gain no benefit from changing the order of elements from the order in which they appear in the exported file, and, in some cases, you might invalidate the XML structure by changing element sequence.</p>

Every XML tag fits into one of the following XML tag types:

- Start tag-Defines the beginning of an element. The start tag consists of an open angle bracket (<), a name, zero or more attributes, and a close angle bracket (>). Every start tag must be followed by an end tag at some point in the document.
- End tag-Defines the end of an element. The end tag consists of an open angle bracket and a forward slash (</), followed by the same name defined in its corresponding start tag, and ends with a close angle bracket (>).
- Empty tag-The empty tag is denoted in two forms. If a tag pair has no data between them, the tag pair is considered an empty tag. Officially, according to the XML specification, an empty tag looks something like this:

```
<<empty tag example/>>
```

In this form, the empty tag consists of an open angle bracket (<), followed by an element name, a slash and a close angle bracket (/>). When you see an empty tag in your configuration files, it signifies an element that the schema requires to be included in the XML file, but whose data is optional.

Start tags can contain attributes, and tag pairs (elements) can contain additional elements. The following example shows an XML file for the Users object. In this example, you see only the Administrator configuration settings.

```
<configuration xmlns="http://xml.juniper.net/ive-sa/6.2R1"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <authentication>
    <auth-servers>
      <auth-server>
        <local>
          <users>
            <user>
              <username>admin</username>
              <fullname>Platform Administrator</fullname>
              <password-encrypted>3u+U</password-encrypted>
              <one-time-use>>false</one-time-use>
              <enabled>>true</enabled>
              <change-password-at-signin>>false</change-password-at-signin>
            </user>
          </users>
        </local>
        <name>Administrators</name>
      </auth-server>
```

You make changes to the string data that is displayed between start and end tags. For example, using the preceding example, you can add to or change the following elements:

- <username>**admin**</username>
- <fullname>**Platform Administrator**</fullname>
- <password-cleartext>**password**</password-cleartext>
- <change-password-at-signin>**false**</change-password-at-signin>
- <name>**Administrators**</name>



The preceding sample displays the password element's data as encrypted data. You can modify the password if you change the element to password-cleartext. If you modify the password, the password value is visible until it is imported back into the system. Once imported, the system encrypts the password.

If you enter passwords for new users in cleartext format, the passwords are visible in the file, therefore, you might consider setting the Change Password at Next Login option to true.

- Because passwords are encrypted by default, they are fully portable from one system to another.
- Use the password-cleartext element and enter a text password when changing passwords through the XML file.



If you change a user for a given authentication server or an authentication server for a given user, you are creating a different user, not updating an existing user or authentication server. User and authentication server together logically define a unique user.

Comparing Configuration Settings and Values Shown in the User Interface with the Ones in the XML File

The elements in the XML file are closely related to the objects and their options as you see them in the admin console. The element names in the XML instance file correlate closely with the displayed object and option names.

For example, select Users > User Roles > [Role] > General > Session Options. The admin console renders the possible values for a roaming session as an option button group, consisting of the values:

- **Enabled**
- **Limit to subnet**
- **Disabled**

The following snippet from the exported configuration file shows the session options for the Users role. On the bolded line, the roaming session option is disabled:

```
<session-options><SessionOptions>  
  <MaxTimeout>60</MaxTimeout>  
  <RoamingNetmask />  
  <Roaming>disabled</Roaming>  
  <PersistentSession>>false</PersistentSession>  
</SessionOptions>
```

In the schema file, you can locate the allowable values for the roaming session option:

```
<Attribute roaming:START>
```



```

<xsd:element name="roaming" minOccurs="0">
...
  <xsd:enumeration value="enabled">
...
  <xsd:enumeration value="limit-to-subnet">
...
  <xsd:enumeration value="disabled">
...
</xsd:element>
<Attribute roaming:END>

```

To change the value for the roaming session from **Disabled to Limit to subnet**, replace disabled with **limit-to-subnet**.

This example shows that the admin console often provides all of the allowable values, displayed either in an option button group, as check boxes, as list boxes, or as other types of user interface components. The XML file displays only the current state of your configuration. The schema file displays all of the actual values for the configuration options that are supported.


Understanding Referential Integrity Constraints

The system configuration objects are part of a data model that is enforced through the use of referential integrity constraints. You cannot change these constraints, but you should understand them before you attempt to delete objects that maintain dependencies to other objects.

If you violate the referential integrity constraints, your import operation fails.

In the following figure the boxes represent object types and the arrows represent dependent relationships between the object types. Arrows point from dependent objects to objects.

Object Referential Integrity Constraints



The system does not allow you to delete an object on which another object depends. Conversely, when you add an object, you must add any other objects on which that object depends.

Sign-in URLs depend upon realms and sign-in pages. Realms depend upon both authentication servers and roles. Policies depend upon roles. Users depend upon authentication servers.

Consider the following scenarios based on the preceding figure:

- If you add sign-in URLs, you must add realms, sign-in pages, roles, and authentication servers. You need to add an authentication server and at least one role to support the realm, and you must add the realm and the sign-in page to support the new sign-in URL.
- If you add a user, you must be able to assign it to an authentication server. If there is no authentication server on the target node yet, you must add one in the XML file.
- If you add a policy, you must be able to assign it to a role. If there is no role on the target system, you must add one in the XML file.
- If you delete an authentication server, you might strand realms and users, therefore, you need to make sure no realms or users depend on the authentication server before you attempt to delete it.
- If you delete a role, you might strand policies and realms. To delete a role, you must first delete any policy that depends upon the role, or reassign associated policies to another role. Also, to delete a role, you must first delete or reassign any realm that depends upon that role.
- If you delete a sign-in page, you might strand one or more sign-in URLs. To delete a sign-in page, you must first delete any associated sign-in URLs or reassign them to other sign-in pages.

Referential integrity checks are performed only during XML import.

Using Operation Attributes

Operation attributes define the positioning or action of XML data within the schema. If you do not specify an operation attribute, the modified data is merged by default.

XML data with an operation attribute has the following format:

```
<object1 xc:operation="operator for object1 and its children unless new operator is defined">
  ...
  <object2>
    ...
    <object3 xc:operation="operator for object3">
      ...
    </object3>
  ...
</object2>
...
```

```
</object1>
```

The operation attribute is applied to all children objects unless a different operation attribute is defined in children objects.

The following operation attributes are supported:

- Merge-The configuration data identified by the element that contains this attribute is merged with the configuration at the corresponding level in the configuration datastore identified by the target parameter. This is the default behavior.
- Replace-The configuration data identified by the element that contains this attribute replaces any related configuration in the configuration datastore identified by the target parameter. Only the configuration actually present in the configuration parameter is affected.
- Create-The configuration data identified by the element that contains this attribute is added to the configuration if and only if the configuration data does not already exist on the device.
- Delete-The configuration data identified by the element that contains this attribute is deleted in the configuration datastore identified by the target parameter.
- Insert before-Changes the position of a configuration element in an ordered set.
- Insert after-Changes the position of a configuration element in an ordered set.
- Rename-Changes the name of one or more of a configuration object's identifiers.

If you are merging a list of objects to an existing list of objects in the configuration store, the results of the merged list might be unexpected. During a merge operation the order of the objects in the new list is not maintained. If you are importing a list of objects and would like to preserve the order of the new list, you should use the replace operation attribute. You can also use insert before or insert after to ensure that you produce the hierarchy that you intended.

Operation attributes are applied to elements recursively unless new operators are also defined within lower-level elements. There are limitations on the legal operator that can be used in child elements without conflict with the parent operator. The following table displays the legal operator relationships between parent and child elements.

The following table lists the Legal Operator Attribute Relationships:

Child >							
V-Parent	Create	Merge	Replace	Delete	Insert		
before	Insert						

Child >							
after	Rename						
None	OK	OK	OK	OK	OK	OK	OK
Create	OK	OK	Error	Error	OK	OK	Error
Merge	OK	OK	OK	OK	OK	OK	OK
Replace	Error	OK	OK	Error	OK	OK	Error
Delete	Error	OK	Error	OK	Error	Error	Error
Insert							
before	OK	OK	OK	OK	OK	OK	OK
Insert							
after	OK	OK	OK	OK	OK	OK	OK
Rename	OK	OK	OK	OK	OK	OK	OK

The following examples demonstrate the import operation:

Example 1: Set the MTU to 1500 on an interface named "Ethernet0/0" in the running configuration.

```
<interface>
  <name>Ethernet0/0</name>
  <mtu>1500</mtu>
</interface>
```

Example 2: Add an interface named "Ethernet0/0" to the running configuration, replacing any previous interface with that name.

```
<interface xc:operation="replace">
  <name>Ethernet0/0</name>
  <mtu>1500</mtu>
  <address>
    <name>192.0.2.4</name>
    <prefix-length>24</prefix-length>
  </address>
</interface>
```

The default import modes have the following equivalent attributes on the root object of the configuration tree:



- Standard Import is always a merge operation.
 - Quick Import is a create operation.
 - Full Import is a replace operation.
-

Using the Push Configuration Feature

This topic describes the push configuration feature.

Push Configuration Overview

The push configuration feature supports simple configuration management across an enterprise without requiring you to deploy the systems as a cluster. You push a partial configuration from the running configuration on the source system to the running configuration on one or more target systems.

It is not desirable to push some groups of settings to a running configuration, so the following groups of settings are not supported:

- Network configurations
- Licenses
- Cluster configurations
- Certificates
- SNMP settings
- Syslog server settings
- Push configuration targets

Guidelines and Limitations

The following table lists the Push Configuration Guidelines and Limitations:

Category	Guidelines and Limitations
General	<p>The following guidelines and limitations apply:</p> <p>You can push a configuration to systems running the same software version (same build number) or higher software version.</p> <p>The source device pushes data over the management port (if configured) or the internal port. The target device can receive data over the internal or external port or management port.</p> <p>You can push to a single target or to multiple targets. For example, if you install several new systems, you can push a common configuration to set their initial configuration.</p> <p>When a configuration push job begins on a target, no warning is displayed, and the administrators are automatically logged out to avoid potential conflicts.</p> <p>For selected configuration push, if the configuration to be pushed contains one or more Pulse client packages, we recommend you push the Pulse client packages first considering one Pulse client package per push and then push the remaining configurations.</p> <p>When the job has completed on a target, the target device restarts its services. Brief interrupts might occur while the service restarts. You must push to targets when they are idle or when you can accommodate brief interruptions.</p> <p>Immediately after pushing an Active Directory authentication server configuration, you must edit the configuration to change the Computer Object name. Unexpected problems might arise if two systems join an Active Directory domain using the same Computer Object name.</p> <p>You must delete the failed push jobs before performing a new push.</p> <p>When performing an entire configuration push, with changes to settings such as security settings, the target web server might get restarted during configuration import. As a result, the source might experience a lost connection. You can resume the job from the source at a later point to see the result of the import.</p>
Licenses	The push configuration job does not push licenses or licensing settings.

Category	Guidelines and Limitations
Clusters	<p>You can push a configuration to multiple targets, as long as targets are not part of the same cluster.</p> <p>You must not perform the clustering operations such as adding a cluster, deleting a cluster, and so on when performing a push configuration. If such events occur, then unsuccessful jobs will be aborted, and the backup files will be deleted.</p> <p>You can push a configuration to multiple targets, as long as targets are not part of the same cluster.</p> <p>You must not use VIPs during push configuration. Instead you must use the internal IP or the management IP of one of the nodes to create the target.</p> <p>You must delete the backed up configuration on the target node(s) as soon as possible to free up the disk space.</p>

Configuring Targets

To configure push configuration targets:

1. Select **Maintenance > Push Config > Targets** to display the target list and source options configuration page.


The Push Configuration Target List and Source Device Settings Page - Ivanti Connect Secure figure shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration for the source options as described in [Table](#).
3. Click **New Target** to display the configuration page for targets.

[Figure](#) shows the configuration page for Ivanti Connect Secure.

4. Complete the configuration as described in [Table](#).
5. Save the configuration.

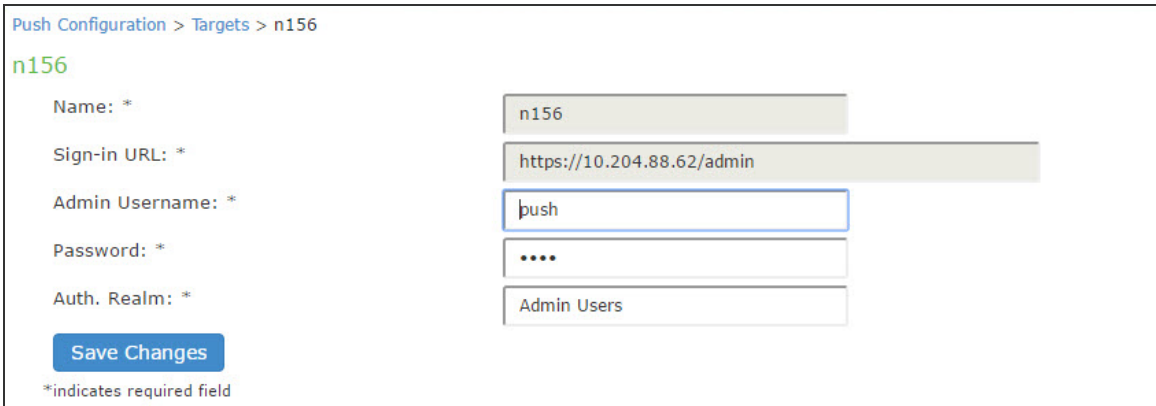
Push Configuration Target List and Source Device Settings Page - Ivanti Connect Secure:



The following table lists the Push Configuration Target Source Device Configuration Options:

Settings	Guidelines
Allow this system to be a target	Select this option to allow the system to accept configuration pushed from another system. This option must be selected on targets, but does not have to be selected on the source system.
Validate target server certificate	Select this option on the source system if you want the source system to validate the target system server certificate before pushing the configuration.
Save Changes	Click this button if you have changed the source device configuration options described above.
Delete	Select a row in the table and click Delete to remove the target from the list. You cannot delete a target if it has push configuration results associated with that target.

Push Configuration Targets Configuration Page:



The following table lists the Push Configuration Targets Configuration Guidelines:

Settings	Guidelines
Name	Specify a name to identify the target within the system. Target names cannot be edited after they have been saved.
Sign-in URL	Specify the URL for the administrator sign-in page. Sign-in URLs cannot be edited after they have been saved.
Admin Username	Specify an account on the target system that the push configuration job can use to sign-in and make changes to the configuration. The job can make wide-ranging configuration changes, so the user must have full administrative privileges. In other words, the user must belong to the .Administrators role.
Password	Specify the corresponding password.
Auth. Realm	Specify the administrator authentication realm on the target system. The access management framework must be configured so that the job process (run as the username specified above) can sign in without any human interaction. For example, you cannot have dynamic credentials or multiple roles that are not merged, as these both require manual interaction. We recommend that you create an administrator account on each target that can be used exclusively for push configuration. Configure the administrator realm so that the realm policy and role mapping rules do not result in prompts requiring human interaction. For example, the user must be able to log in with static password authentication or two-factor tokens that do not use challenge-response type authentication. For example, certificates, Soft ID, and Defender Authentication are not supported.

Configuring Push Settings

To configure the settings to be pushed:

1. Select **Maintenance > Push Config > Push Configuration** to display the configuration page.

The following Push Configuration Selected Settings Page figure shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration and push configuration operation as described in Table Push Configuration Selected Settings and Action Guidelines.

Push Configuration Selected Settings Page:

Push Configuration > Push Configuration

Push Configuration

- Push Configuration**
- Targets
- Results
- History

What to push: Selected configuration ▾

Push selected configuration to other device(s).

- Expand All
- Select All

- ▶ **System Settings...** none selected
- ▶ **Sign-in Settings...** none selected
- ▶ **Endpoint Security...** none selected
- ▶ **Authentication Realms...** none selected
- ▶ **Roles...** none selected
- ▶ **Resource Profiles...** none selected
- ▶ **Resource Policies...** none selected
- ▶ **Pulse Secure client...** none selected
- ▶ **Enterprise Onboarding...** none selected
- ▶ **Local User Accounts...** none selected
- ▶ **Maintenance Settings...** none selected

▼ Push configuration

Available Targets:

n156

Selected Targets:

(none)

- Add ->
- Remove

Copyright © 2022, Ivanti, Inc. All Rights Reserved. [Privacy and Legal](#).

- Overwrite duplicate settings
- Allow Rollback to Previous Config

Push Configuration Selected Settings Page:

Push Configuration > Push Configuration

Push Configuration

Push Configuration | Targets | Results | History

What to push: Entire configuration

Push this device's configuration to other device(s).

Push configuration

Available Targets: n156

Selected Targets: (none)

Add -> Remove

Overwrite duplicate settings

Allow Rollback to Previous Config

If rollback is enabled, we will take a backup of the full config on all selected targets. The backup file is counted against the limit for push config disk usage. If the complete.

Description:

Schedule Import On Target:

Date: (mm/dd/yyyy)

Time: AM (hh:mm) (GMT-08:00) Pacific Time (US & Canada); Tijuana

Push Configuration

The following table lists the Push Configuration Selected Settings and Action Guidelines:

Settings	Guidelines
Select Settings and Export	

Settings	Guidelines
What to push	<p>Select Selected configuration or Entire configuration.</p> <p>If you select Selected configuration, the page displays controls to select settings groups.</p> <p>If you select Entire configuration, all settings from the source system are pushed, except for the following:</p> <ul style="list-style-type: none"> Network configurations Licenses Cluster configurations Certificates SNMP settings Syslog server settings Push configuration targets
Expand All	Click this button to expand the display of all settings for all groups.
Select All	Click this button to select all settings for all groups.
Settings	
System	<p>Expand this group and select settings found under the System menu.</p> <p>You cannot push host-specific network settings to a target. If you want to copy these settings to another system, use the configuration XML file import/export feature.</p>
Sign-in	Expand this group and select settings found under the Sign-in menu.
Endpoint Security	<p>Expand this group and select settings found under the Endpoint Security menu.</p> <p>ESAP packages are encrypted when exported.</p>
Authentication Realms	Expand this group and select authentication realm settings, including user realms and MAC address authentication realms.
Roles	Expand this group and select settings found under the Roles menu.
Resource Profiles	<p>Ivanti Connect Secure only.</p> <p>Expand this group and select settings resource profiles settings.</p>
Resource Policies	Expand this group and select settings resource policies settings.
Pulse Client	Expand this group and select settings found under the Ivanti menu.
Local User Accounts	Expand this group and select local authentication server settings.
Maintenance	Expand this group and select settings found under the Maintenance menu.

Settings	Guidelines
Push Configuration	
Available Targets / Selected Targets	Use the Add and Remove buttons to select the targets.
Overwrite duplicate settings	<p>Select this option to overwrite settings on the target that have the same name as settings being pushed.</p> <p>If you do not select this option, the push configuration job copies only configuration objects that have names different from the configuration objects on the target.</p>
Allow Rollback to previous configuration	<p>Select this option to revert to a previous configuration state, effectively rolling back configuration changes.</p> <p>If you select this option, the local configurations on the target node will be backed up before importing the configurations. You can also undo the push configuration if you want to discard the changes and revert back to the previous state. We recommend you delete the backed-up configuration if the import is successful.</p> <p>If the target configuration is large the rollback of configurations can take several minutes to complete.</p>
Description	Enter the description for the job. The job description is limited to 100 characters.
Schedule Import on Target	<p>Select this option to allow a delayed import on the target node. If you select this option, the selection applies to all the targets in the job. The import schedule is measured in HH:MM (hours, minutes) format. The schedule is specified according to source's timezone.</p>
Push Configuration	<p>Click this button to push the selected configuration data to the specified targets.</p> <p>You can pause the push for a target during the push process. If errors occur during the push process, the job stops, and the configuration for the target is not imported. However, you can resume the failed push jobs. Error messages are displayed on the Results page.</p> <p>If you have specified multiple targets and a push configuration job to a target fails, the job continues to the next target until specified targets are updated (or fail). The results page displays the status and any problems encountered during the process.</p>

Viewing Configuration Push Results

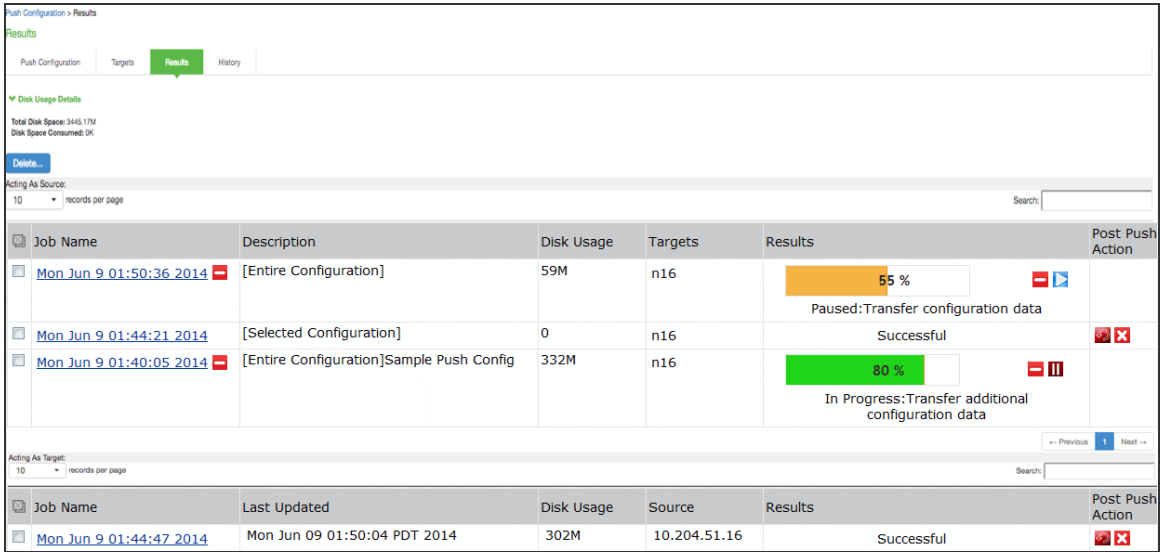
Purpose	The source system saves and displays the push configuration results in the Results tab.
Action	To view push configuration job results:

1. Select **Maintenance > Push Config > Results** to display the results page.

The Push Configuration Results Page figure shows the results page for Ivanti Connect Secure. The push configuration results page auto refreshes for every 30 seconds.

2. Examine the results to verify success or learn the reasons the push job failed.
3. Click the job name to display additional information about the job.
4. Select a job and click Delete to remove it from the results page.

Push Configuration Results Page



The following table describes the information displayed on the Results page and the various management tasks you can perform.

The following table lists the Push Configuration Results:

GUI Element	Guidelines
Disk Usage Details	<p>Displays the disk space available for push configuration and the disk space consumed by all the push jobs in the device.</p> <p>The disk space consumed by individual push jobs are also mentioned across each push job under the disk usage column. When total disk space consumed reaches the total disk space push jobs may fail and you can see the results column to see the failure message. You need to monitor the disk space consumed by push configuration to avoid push failures related to disk space limits.</p>
Description Column	Displays the type of the push configuration.
Disk Usage Column	Displays the disk space used by the job.
Results Column	<p>Displays the status of the transfer and result of post push action. It also displays the status of the push such as login, export, transfer, backup, import and so on. The status result message shows the type of data that is getting transferred. For a paused or failed target, the information on the current state of the job when it is paused, or failure reasons if any is displayed. This column also shows the progress of data transfer using a bar chart. For selected push additional configuration data (additional configuration data refers to configuration that is transferred only if it is modified or not available on the target) includes ESAP package, Pulse client package, VDI configurations, Terminal services, Host Checker files, Custom sign in pages and notifications, and Applet files. For complete configuration push additional data includes ESAP and Pulse client packages.</p>
Post Push Action	<p>Displays the options that the user can perform after the push such as roll back and delete backup. It also displays the post push actions such as rollback done, backup deleted, rollback failed, performing rollback, deleting back up and so on.</p>
Resume	Select this option to resume a paused or a failed push.
Undo	<p>Select this option to rollback to previous configuration that was backed up. Note that you can perform this operation only when the push is successful and Allow Rollback to Previous Configuration is selected. This option is available only if the backup is not deleted or undo is not done yet.</p>
Abort	Select this option to cancel an entire push job or push to particular target within a job. An aborted push cannot be resumed.

GUI Element	Guidelines
Pause	Select this option to temporarily pause the push operation to a specified target.
Delete Backup	Select this option to delete the backup configuration on the specified target. Note that this option is available only when the users selects the Allow Rollback to Previous Configuration option during the push job.

Viewing Configuration Push History

Purpose	The source/target system saves and displays up to 5 push history results per target/source in the History tab. When the history table reaches 5 entries, the system removes the oldest result data when the next push configuration job is started.
Action	To view push configuration push history:

1. Select **Maintenance > Push Config > History** to display the history page.

The following figure shows the history page for Ivanti Connect Secure.

2. Examine the history to verify success or learn the reasons the push job failed. The history page displays rollback history however the failure reason is not displayed. You can check the failure reason in the details page for each job. It also displays the timestamp history information of successful, failed push jobs, or if a configuration is undone.
3. Select the source name/target name and click Delete History to remove it from the History page.

Push Configuration History Page

Push Configuration > History

History

Push Configuration Targets Results History

Delete History

10 records per page Search: []

Source Name	Source IP	Source History
<input type="checkbox"/> localhost2	10.204.51.16	Push Successful: Mon Jun 9 01:50:04 2014
<input type="checkbox"/> localhost2	10.204.51.25	Push Successful: Thu May 29 23:43:56 2014
		Push Successful: Thu May 29 23:16:56 2014 Undone: Thu May 29 23:19:39 2014
		Push Successful: Thu May 29 20:35:15 2014

10 records per page Search: []

Previous 1 Next

Target Name	Target Sign-in URL	Target History
<input type="checkbox"/> n16	https://10.204.51.16/admin	Push Successful: Mon Jun 9 01:45:53 2014
		Push Successful: Thu Jun 5 12:12:21 2014
		Push Failed: Thu Jun 5 12:10:05 2014
		Push Successful: Thu Jun 5 11:59:00 2014
		Push Failed: Thu Jun 5 11:51:02 2014
<input type="checkbox"/> n156	https://10.204.50.156/admin	Push Successful: Fri Jun 6 00:18:47 2014
		Push Failed: Fri Jun 6 00:16:04 2014
		Push Failed: Thu Jun 5 00:30:50 2014

System Maintenance

Using the System Maintenance Pages

You can use the System > Maintenance pages to perform the following tasks:

- Enable system maintenance options, such as software version monitoring and disk clean-up.
- Upgrade, downgrade, or rollback the system software.
- Download client installer files so that you can distribute them in out-of-band methods to end users.
- Test network connectivity between the system and servers that have been configured to be used with it.
- Display hardware status.

Configuring System Maintenance Options

You can use the maintenance options page to enable various system maintenance features.

To enable various system maintenance features:

1. Select **Maintenance > System > Options** to display the maintenance options page.
2. Select options as described in the following table.
3. Save the configuration.

The following table lists the System Maintenance Options Configuration GuidelinesE:

Options	Guidelines
Automatic version monitoring	<p>If you enable this option, the system reports to Ivanti the following data:</p> <ul style="list-style-type: none"> Machine identifier. Information describing your current software, including: <ul style="list-style-type: none"> Software build number and build name. An MD5 hash of your license settings. An MD5 hash of the internal interface IP address. If this node is in a cluster, the number of nodes within that cluster. Current state of the node. Cluster type (active/active, active/passive). Total number of unique subnets on the cluster nodes. Version of Pulse client. Version of ESAP. Cluster log synchronization status. Total number of concurrent users on the device. Number of Ivanti tunnels. <p>We strongly recommend that you enable this service.</p>
Gzip compression	<p>Connect Secure only. Use gzip compression to reduce the amount of data sent to browsers that support HTTP compression. This can result in faster page downloads for some users.</p>
Kernel Watchdog	<p>Enables the kernel watchdog that automatically restarts the system under kernel deadlock or when kernel runs low on some key resources.</p> <p>Enable the kernel watchdog only when instructed by Ivanti Technical Support.</p>
Resource throttling	<p>Enables system resource throttling in the system that gives system processes higher priority. High priority processes will get high resources under system load. Changing this option will cause a system reboot.</p>
File System Auto-clean	<p>Enables the system to automatically clean up the file system when disk utilization reaches 90%.</p> <p>The clean-up operation deletes files that might be relevant in debugging- for example, debug logs, core files, and snapshots might be deleted.</p>

Options	Guidelines
Web installation and automatic upgrade of Pulse Clients	<p>After you deploy Pulse client software to endpoints, software updates occur automatically. A Pulse client can receive updates from the server. If you upgrade the Ivanti software on your Ivanti server, updated software components are pushed to a client the next time it connects.</p> <p>A bound endpoint receives connection set options and connections from its binding server, but it can have its Pulse client software upgraded from any Ivanti server that has the automatic upgrade option enabled. During a client software upgrade the client loses connectivity temporarily.</p>
Enable Pulse Client Components removal Tool for Cert issue Remediation	<p>Provides an option for the Admin to enable users to download the Pulse Client Components removal (Pulse Upgrade Helper) tool on Windows End User machines upon Browser access and remediates the certificate expiry issue. For more information, refer KB44781 and KB44810.</p>
Virtual Terminal console	<p>Enables the virtual terminal on a virtual appliance. Clear this check box to use the serial console. Changing this setting will restart the system.</p>
Java instrumentation caching	<p>Connect Secure only. Caches the Java instrumentation to improve the performance of Java applications.</p>
Show Auto-allow	<p>Connect Secure only. The auto-allow option provides the means to automatically add bookmarks for a given role to an access control policy, for example, Web bookmarks with auto-allow set are added to the Web access control policy. You only use this feature if you also use Resource Policies. We recommend that you use Resource Profiles instead.</p>
Do not show Task Guidance/Help page on admin login	<p>This option is applicable only in case there are no licenses installed. When enabled, Task Guidance/Help page does not appear automatically upon administrator login.</p>
Clear all configuration data at this device	<p>This option clears all keys and triggers a configuration reset and reboots the device.</p>
Prevent system overload	<p>Disallows user login, user login via Pulse Desktop, HTML5 connection or connection to a web resource when the CPU load is above a certain threshold. By default, this option is disabled for ICS upgrades and enabled for new installation.</p> <p>Exception: Admin logins, DMI and inbound REST calls are not blocked due to CPU overload.</p> <p>When a login to the HTML5 connection or connection to a web resource is blocked and when a user tries to log in, the login page will display an appropriate system busy message.</p>

Options	Guidelines
	<p>To configure log events for User Access, in the System > Log/Monitoring > User Access > Settings tab, select the System Too Busy check box. By default, this option is enabled.</p> <p>Select System > Log Monitoring > User Access > Log to view the logs.</p>
Auto reboot the system	This option automatically reboots the system when the appliance is in kernel panic state.
End-user Localization	<p>Select one of the following options:</p> <ul style="list-style-type: none"> Automatic (based on browser settings) English (U.S.) Chinese (Simplified) Chinese (Traditional) French German Japanese Korean Spanish Italian Polish
External User Records Management	
Persistent user records limit	<p>Specify the maximum number of user records.</p> <p>This feature is useful when system performance is affected due to a large number of user records. We highly recommend you consult Technical Support prior to using this feature. Deleting a user record removes all persistent cookies, SSO information, and other resources for that user. It does not remove the user record from the external or internal authentication server. If you delete a user record and that user logs back in to the authentication server, new user records are created. Records are not removed if that user is currently logged in.</p>
Number of records to delete when the limit is exceeded	Specify a number. Older records are removed first. A user record is not deleted if that user is currently logged in.
Delete records now	Check whether the persistent user records limit has been exceeded. If it is, delete the number of user records specified in the option above.

Options	Guidelines
Automatic deletion of user records periodically	Check whether the persistent user records limit will be exceeded whenever a new user record is about to be created. If true, delete the records prior to creating the user new record.

Upgrading the System Software

This topic describes how to upgrade, downgrade, and rollback the system software.

Downloading a Software Package

To download a software package:

1. Go to <https://www.pulsesecure.net/support/> and browse to the software download page for your product.
2. When prompted, log in with your Ivanti customer username and password.
3. Accept the license agreement.
4. When prompted, save the software package to your local host.

Uploading a Software Package

You can upload a software package to the system without immediately initiating the upgrade process. This is known as staging the upgrade. You can stage one package. Uploading a second package overwrites the previous staging.

To upload a software package:

1. Select **Maintenance > System > Upgrade/Downgrade** to display the system software maintenance page.

The following figure shows Ivanti Connect Secure.

2. Under **Managed Staged Service Package**, select **Upload new package** into staging area and use the Browse button to locate and select the service package file.
3. Click **Submit** to upload the file.

The Upload Status window shows the progress of the upload operation.

Software Upgrade Page

System Maintenance > Install Service Package

Install Service Package

Platform Upgrade/Downgrade Options Installers

Installing a service package can take several minutes and requires the system to reboot. Because existing system data is backed up during this process, you can decrease installation time by clearing your [system log](#) before trying to install a service package.

Note: Browsing away from this page while uploading the package will abort the installation.

▼ Install Service Package

From File

No file chosen

From Staged Package

Choose this option if you want to install the staged service package.

DELETES all system and user configuration data before installing the service package, restoring the member to an unconfigured state. Use this option if you want to downgrade to an older service package than the currently installed package. **Do NOT check this box** if you want to retain existing settings and data during a system upgrade to a newer service package.

Note: This option does not change the factory image.

▼ Manage Staged Service Package

Upload new package into staging area

No file chosen

Delete Staged Package



If you have enabled logging for Administrator changes (**System > Log/Monitoring > Admin Access > Settings** page), a log is written to the Admin Access logs page.

Upgrading the System Software

Installing a service package can take several minutes and requires the system to reboot. Because existing system data is backed up during this process, you can decrease installation time by clearing your system log before trying to install a service package.

- When the system software is upgraded to 9.0R4:
 - latest set of **Trusted Server CAs** are uploaded. These new set of **Trusted Server CAs** will be seen in the **System > Configuration > Certificates > Trusted Server CAs** page.
 - Any expired certificates in the default Trusted Server CA store are removed from the system.
- When the system software is upgraded to 9.0R3, it automatically upgrades Ivanti Connect Secure to OpenSSL version 1.0.2n.

To upgrade the operating system:

1. Select **Maintenance > System > Upgrade/Downgrade** to display the system software maintenance page.

[Software Upgrade Status Page](#) shows the system software maintenance page.

2. Under Install Service Package, select one of the following options to proceed:
 - **From File**-Use the **Browse** button to locate and select the service package file.
 - **From Staged Package**-Select the service package file that was previously uploaded.



Do not select the Deletes option when you are upgrading software. The Deletes option is available to support downgrading software.

3. Click **Install**.

The system displays the Service Package Installation Status page, which provides a summary of the integrity checks and compatibility checks and other status indicators.

Software Upgrade Status Page

Service Package Installation Status

The installation process takes a few minutes. When complete, the system needs to reboot. Please wait...

- Step 1: Verifying package integrity complete (37 seconds)
- Step 2: Extracting install script complete (114 seconds)
- Step 3: Extracting install script complete (19 seconds)
- Step 4: Running system compatibility checks ... complete (0 seconds)
- Step 5: Saving copy of system config complete (59 seconds)
- Step 6: Boot partition is set to xda device complete (1 seconds)
- Step 7: Preparing disk partitions complete (9 seconds)
- Step 8: Extracting contents of new package complete (35 seconds)
- Step 9: Saving package complete (115 seconds)
- Step 10: Finalizing installation ... complete (1 seconds)
- Step 11: Encrypting drive please wait
.....
complete (321 seconds)
- Step 12: Switching current system to "rollback" and enabling new system ... complete (0 seconds)
- Step 13: Boot partition is set to xda device complete (3 seconds)

ⓘ Installation completed successfully and the system will now reboot.: Note that the Administrator Console will be unavailable while the system reboots.(Watch the serial console for messages).
 When the system reboots click [here](#) to continue using the Administrator Console.



If you have enabled logging for Administrator changes (**System > Log/Monitoring > Admin Access > Settings** page), a log is written to the Admin Access logs page. If you have enabled logging for System Status (System > Log/Monitoring > Events > Settings page), logs are written to the Events logs page.

Downgrading the System Software

If necessary, you can downgrade to an earlier version of the system software. When you downgrade, you must clear the system and configuration data to avoid unexpected behavior that can occur when the system has data that relates to the newer software.

If you downgrade the system, you must reestablish network connectivity before you can reconfigure it.

To downgrade the operating system:

1. Select **Maintenance > System > Upgrade/Downgrade** to display the system software maintenance page.
[System Maintenance Platform Page](#) shows the system software maintenance page.
2. Under Install Service Package, select one of the following options to proceed:
 - **From File**-Use the **Browse** button to locate and select the service package file.
 - **From Staged Package**-Select a service package file that was previously uploaded.
3. Select the Deletes option to delete all system and user configuration data before installing the service package, restoring the member to an unconfigured state.
4. Click **Install**.

Rolling Back the System Software

If necessary, you can roll back the system to the previous software version and configuration state. The system is rebooted and unavailable for a few minutes when you roll back.

To roll back the operating system:

1. Select **Maintenance > System Maintenance > Platform** to display the system maintenance platform page.
[System Maintenance Platform Page](#) shows the system maintenance platform page for Ivanti Connect Secure.
2. Click **Rollback**.

System Maintenance Platform Page

System Maintenance > Platform

Platform

Upgrade/Downgrade Change Personality Options Installers

Hostname: localhost2
Model: PSA-5000
Machine ID: 0320MSTR5098B01LS
Serial Number: 0320012018100004
Uptime: 7 days, 16 hours, 6 minutes, 26 seconds
Current version: 9.1R15 (build 17875)
Rollback version: 9.1R13 (build 16675)

Node operations: Restart Services Reboot... Shut Down... Rollback...

Connectivity:
 This will ping various configured servers to test the device's connectivity. Test Connectivity

Hardware Status

Fan Status:

Fan	Status
1	●

Temperature: 40 °C



- The rollback option appears only if you have previously upgraded the system software.
- If you have enabled logging for System Status (System > Log/Monitoring > Events > Settings page), logs are written to the Events logs page.

Downloading Client Installer Files

You can use the system maintenance client installers page to download client installer files. The downloadable files include .exe and .msi files for use installing clients on Windows platforms, and .dmg files for installing clients on Macintosh platforms.

To download client installer files:

1. Select **Maintenance > System Maintenance > Installers** to display the client installer files page.

[System Maintenance Client Installers Page -Ivanti Connect Secure](#) shows the client installer files for Ivanti Connect Secure.













2. Click **Download** to download the file to your local host.

[System Maintenance Client Installers Page -Ivanti Connect Secure](#)

System Maintenance > Installers

Installers

Platform Upgrade/Downgrade Change Personality Options **Installers**

	Pulse Secure Installer Service (.exe) This component simplifies future installation and upgrades of Pulse Secure's client software for users with limited desktop privileges. Use this self-extracting .exe package unless a specific requirement exists for Microsoft Windows Installer (msi) packages. This package can be deployed with limited user privileges if a previous version of the Installer Service is running.	version: 9.1R15 (build 17875)[9.1.15-17875(4606)] [Download]
	Pulse Secure Installer Service (.msi) This component simplifies future installation and upgrades of Pulse Secure's client software for users with limited desktop privileges. Deploy this Microsoft Windows Installer (msi) package if your organization or infrastructure requires msi packages. One such example would be automated software installation using Microsoft Systems Management Server.	version: 9.1R15 (build 17875)[9.1.15-17875(4606)] [Download]
	Host Checker This component verifies security settings on user's workstation.	version: 9.1R15 (build 17875)[9.1.15-17875(4606)] [Download]
	Third-party Integrity Measurement Verifier (IMV) Server This component allows third-party integrity measurement verifiers (IMV's) to be used with the Pulse Secure Gateway	version: 9.1R15 (build 17875)[9.1.15-17875(4606)] [Download]
	Pulse Secure Installer (Linux 64-bit) This setup application installs New Pulse Secure Linux Client. Installing New Client will not import Connection Profiles from Older Version of Client.	CentOS/RHEL platforms : 9.1.15.15241 [Download] Ubuntu/Debian platforms : 9.1.15.15241 [Download]
	Pulse Secure Installer (Windows 32-bit) This setup application installs all Pulse Secure components. Please consult the Pulse Secure Supported Platforms Guide for this version to determine which versions of Windows 32-bit client operating systems are supported.	version: 9.1.15.15241 [Download]
	Pulse Secure Installer (Windows 64-bit) This setup application installs all Pulse Secure components. Please consult the Pulse Secure Supported Platforms Guide for this version to determine which versions of Windows 64-bit client operating systems are supported.	version: 9.1.15.15241 [Download]
	Pulse Secure Installer (ARM 64-bit) This setup application installs all Pulse Secure components. Please consult the Pulse Secure Supported Platforms Guide for this version to determine which versions of ARM 64-bit processor supported.	version: 9.1.15.15241 [Download]
	Pulse Secure Installer (MacOS) This setup application installs all Pulse Secure components. Please consult the Pulse Secure Supported Platforms Guide for this version to determine which versions of MacOS client operating systems are supported.	version: 9.1.15.15241 [Download]
	Pulse Application Launcher Installer This setup application installs all Pulse Application Launcher components.	version: 9.1R15 (build 17875)[9.1.15-17875(4606)] [Download]
	Pulse Application Launcher Installer (MacOS) This setup application installs all Pulse Application Launcher components.	version: 9.1R15 (build 17875)[9.1.15-17875(4606)] [Download]
	Pulse Application Launcher Installer (Linux) This setup application installs all Pulse Application Launcher components.	CentOS/RHEL platforms : 9.1R15 (build 17875)[9.1.15-17875(4606)] [Download] Ubuntu/Debian platforms : 9.1R15 (build 17875)[9.1.15-17875(4606)] [Download]

Restarting, Rebooting, and Shutting Down the System

You can use the admin console to perform restart, reboot, and shut down operations. The following items explain these terms:

- **Restart**-Kills all processes and restarts the system. The system is available again after a few minutes.
- **Reboot**-Power cycles and reboots the system. The system is available again after a few minutes.
- **Shut Down**-Shuts down the system. The system is not available again until the physical power button on the physical device is used to restart the system.



The restart, reboot, and shutdown operations are applied to all enabled members of a cluster. If you do not want to apply the operations to all members of the cluster, use the System > Clustering > Status page to disable members; then perform the restart, reboot, or shut down operation.

To restart, reboot, or shut down the system:

1. Select **Maintenance > System Maintenance > Platform** to display the system maintenance platform page

[System Maintenance Platform Page](#) shows the system maintenance platform page for Ivanti Connect Secure.

2. Click the desired node operation:

- **Restart Services**
- **Reboot**
- **Shut Down**

System Maintenance Platform Page

The screenshot displays the 'System Maintenance > Platform' page. It features a navigation bar with 'Platform', 'Upgrade/Downgrade', 'Options', and 'Installers'. The main content area shows a server icon and the following details:

- Cluster:** BNG_CONNECT_CLUSTER
- Hostname:** NODE_3_3
- Model:** MAG-SM160
- Serial Number:** 0284072011000123
- Chassis Location:** FPC 3
- Uptime:** 4 days, 2 hours, 11 minutes, 12 seconds
- Current version:** 8.2R3:B2 (build 44679)
- Rollback version:** 8.2R1.1 (build 42861)

Under **Node operations:**, there is a 'Reboot this node...' button. Under **Cluster operations:**, there are buttons for 'Restart Services', 'Reboot...', 'Shut Down...', and 'Rollback...'. A note states: 'Cluster operations affect all nodes in the cluster.' Under **Connectivity:**, there is a 'Test Connectivity' button and a note: 'This will ping various configured servers to test the device's connectivity.'

Hardware Status is expanded to show **Fan Status:**

Fan Tray	Status
2	●
3	●

At the bottom, it shows **Temperature: 34 °C**.



If you have enabled logging for Administrator changes (System > Log/Monitoring > Admin Access > Settings page), a log is written to the Admin Access logs page. If you have enabled logging for System Status (System > Log/Monitoring > Events > Settings page), logs are written to the Events logs page.

Testing Network Connectivity

You can use the admin console to test network connectivity to all the servers with which the system is configured to communicate, for example network services or AAA servers.

To test network connectivity:

1. Select **Maintenance > System Maintenance > Platform** to display the system maintenance platform page.

[System Maintenance Platform Page](#) shows the system maintenance platform page for Ivanti Connect Secure.

2. Click **Test Connectivity**.

Server connectivity results are highlighted in the figure.

System Maintenance Platform Page

The screenshot displays the 'System Maintenance > Platform' page. The 'Platform' tab is active, showing system information and hardware status. A red box highlights the 'Server connectivity results' section, which lists five successful ping tests to various destination hosts.

System Maintenance > Platform

Platform | Upgrade/Downgrade | Change Personality | Options | Installers

Hostname: localhost2
Model: PSA-5000
Machine ID: 0320MSTR509880ILS
Serial Number: 0320012018100004
Uptime: 7 days, 16 hours, 10 minutes, 22 seconds
Current version: 9.1R15 (build 17875)
Rollback version: 9.1R13 (build 16675)

Node operations: Restart Services | Reboot... | Shut Down... | Rollback...

Connectivity:
 This will ping various configured servers to test the device's connectivity. **Test Connectivity**

Hardware Status

Fan Status:

Fan	Status
1	●

Temperature: 39 °C

Server connectivity results

Destination host 10.96.112.1 used as Gateway Address is responding.
Destination host 10.64.109.240 used as DNS Server is responding.
Destination host 10.64.104.121 used as Radius Server is responding.
Destination host 10.64.109.189 used as NIS Server is responding.
Destination host 10.64.109.240 used as LDAP Server is responding.

Logging and Monitoring

Logging Overview

The system generates event logs related to system performance, administrator actions, network communications, access management framework results, user sessions, and so forth. The system supports the following log collection methods:

- Local log collector and log viewer.
- Reporting to syslog servers.
- Reporting to SNMP servers.

The following table describes the event log severity levels.

The following table lists the Event Log Severity Levels:

Severity Level	Description
Critical (level 10)	The system cannot serve user and administrator requests or loses functionality to a majority of subsystems.
Major (levels 8-9)	The system loses functionality in one or more subsystems, but users can still access the system for other access mechanisms.
Minor (levels 5-7)	The system encounters an error that does not correspond to a major failure in a subsystem. Minor events generally correspond to individual request failures.
Info (levels 1-4)	The system writes an informational event to the log when a user makes a request or when an administrator makes a modification.

In addition to managing system logs, you can use the admin console to configure collection of client-side logs, including:

- Host checker
- Windows Secure Application Manager
- Java Secure Application Manager and Applet Rewriting
- VPN Tunneling
- Terminal Services

- Virtual Desktops

Configuring Events to Log

To configure log event categories:

1. Select **System > Log/Monitoring**.
2. Click the **Settings** tab to display the configuration page.
[Log Events Settings Configuration Page](#) shows the configuration page.
3. Complete the configuration as described in [Table](#).
4. Save the configuration.



To configure log events for each local log category, you must perform this procedure on each local log tab: Events, User Access, and Admin Access.

Log Events Settings Configuration Page

The screenshot displays the 'Log settings' configuration page. It features a navigation bar with tabs for 'Events', 'User Access', 'Admin Access', 'Sensors', 'Client Logs', 'SNMP', and 'Statistics'. The 'Events' tab is active. Below the navigation bar, there are 'Save Changes' and 'Reset' buttons. The 'Maximum Log Size' section shows a text input for 'Max Log Size' set to '200' MB, with a note to refer to the 'Archiving' page for archiving log data. The 'Select Events to Log' section contains a grid of checkboxes for various event categories, including Connection Requests, System Status, Rewrite, System Errors, MDM API Trace, Pulse One Events, HTML5 Access Events, Statistics, Performance, Reverse Proxy, and Email Proxy Events. The 'Syslog Servers' section includes a note about logging events locally or to external servers, a 'Delete' button, and a table for adding servers. The table has columns for 'Server name/IP', 'Facility', 'Type', 'Client Certificate', and 'Filter'. A row is shown with 'LOCAL0' for Facility, 'UDP' for Type, and 'Standard: Standard (default)' for Filter. An 'Add' button is at the end of the row.

Server name/IP	Facility	Type	Client Certificate	Filter
	LOCAL0	UDP	Select Client Cert	Standard: Standard (default)

The following table lists the Log Events Settings:

Settings	Guidelines
Maximum Log Size	
Max Log Size	<p>Specify the maximum size of the local log. The default is 200 MB. The maximum is 500 MB. The default is a good choice for logs formatted with the Standard format. If you use a more verbose format, such as WELF, specify a larger value.</p> <p>When the local log reaches the maximum log size, the current data is rolled over to a backup log file. A new, empty, file is then created for all subsequent (new) log messages. The log viewer displays the most recent 5000 log messages (the display limit). If the current log file contains fewer than 5000 log messages, older log messages from the backup log file can be displayed, up to a total of 5000 log messages. This makes the log files appear as one, even though they are stored separately.</p> <p>When you save the log messages or use the FTP archive function, the backup log file is appended to the current log file and is then downloaded as one log file. If the log files are not archived or saved by the time they are rolled over again, the oldest log messages (saved in the backup log file) are lost.</p>
Archiving	Click the Archiving link to display the configuration page for Archiving jobs, including log archiving.
Select Events to Log - Events Tab	
Connection Requests	Log events related to connection requests.
System Status	Log events related to changes in system status.
Rewrite	Log events related to rewrite policies.
System Errors	Log events related to system errors.
Statistics	Log user access statistics reported on the System > Log/Monitoring > Statistics tab. If you unselect the Statistics option, the statistics are not written to the log file, but are still reported on the statistics page.
License Protocol Events	Log events related to licensing.
Reverse Proxy	Logs events related to reverse proxy information.
Select Events to Log - User Access Tab	
Login/logout	Log events related to sign in and sign out.

Settings	Guidelines
SAM/Java	Log events related to user access to SAM/Java in the local log file.
User Settings	Log events related to changes to user settings in the local log file.
Client Certificate	Log events related to certificate security.
IF-MAP Client User Messages	Log events related to IF-MAP.
Pulse Client Messages	Log events related to Pulse clients.
HTML5 Access	Log events related to HTML5 access.
Web Requests	Log events related to user access to web.
File Requests	Log events related to user access to files.
Secure Terminal	Log events related to user access to secure terminal.
VPN Tunneling	Log events related to user access to VPN tunneling.
SAML	Log events related to user access to SAML.
System Too Busy	Log events related to ICS overload.
Unauthenticated Web Requests	Log events related to web requests before authentication. By default, this checkbox is disabled.
Select Events to Log - Admin Access Tab	
Administrator changes	Log events related to configuration changes.
Administrator logins	Log events related to administrator access.
License changes	Log events related to licensing.

Enabling Client-Side Logging

Client-side logging is not enabled by default. If necessary, you can enable client-side logging to troubleshoot any client application issues.

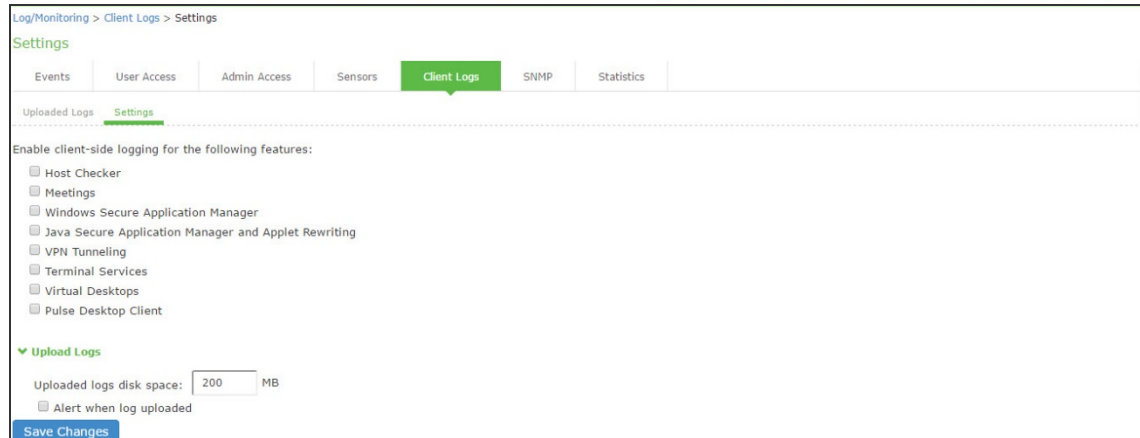
To enable client-side logging:

1. Select **System > Log/Monitoring**.

Click the **Client Logs** tab to display the configuration page. Figure 244 shows the configuration page for Ivanti Connect Secure. Complete the configuration as described in [Table](#).

2. Save the configuration.

Client Logs Configuration Page



The following table lists the Client-Side Logs Settings:

Settings	Guidelines
Host Checker	Select this option to enable client-side logging of Host Checker.
Windows Secure Application Manager	Select this option to enable client-side logging of PSAM.
Java Secure Application Manager and Applet Rewriting	Select this option to enable client-side logging of JSAM and applet.
VPN Tunneling	Select this option to enable client-side logging of VPN tunneling.
Terminal Services	Select this option to enable client-side logging of terminal services.
Virtual Desktops	Select this option to enable client-side logging of virtual desktops.
Pulse Desktop Client	Select this option to enable client-side logging of Pulse desktop clients.

Settings	Guidelines
Upload logs	
Upload logs disk space (MB)	Specify the amount of disk space (in Megabytes) you want to allocate for uploaded client log files. You can allocate disk space from 0 to 200 MB.
Alert when log uploaded	Select this option to receive an alert message when an end user pushes a log file.

Enabling and Viewing Client-Side Log Uploads

If you enable client-side logging for system features, you can also enable automatic upload of those logs at the role level. When you do, end users who are members of the enabled roles can choose to push their log files up to the system at will. Then, you can view the uploaded files through the System > Log/Monitoring > Client Logs > Uploaded Logs page of the admin console.

When you upload log files to a device that is a node in a cluster, keep the following guidelines in mind:

- You can use the Log Node column on the System > Log/Monitoring > Client Logs > Uploaded Logs tab to view the location of existing log files collected by nodes in the cluster. This is specific to a cluster setup and does not apply to a single deployment.
- The user uploads logs to the cluster node to which he is connected.
- You can view upload log entries across all nodes in a cluster. You can save and unzip your uploaded log files from the respective nodes in the cluster where the user uploaded the logs.
- When a node is removed from a cluster, the system deletes the logs of that node from the Uploaded Log List in the cluster and from the node.

To enable end users to upload logs to the system:

Select **Users > User Roles > Select Role > General > Session Options**.

1. In the Upload logs section, select the **Enable Upload Logs** check box.
2. Click **Save Changes**.

Viewing Uploaded Client-Side Logs

If you enable end users to push log files up to the system, you can view the uploaded logs through the System > Log/Monitoring > Client Logs > Uploaded Logs page of the admin console. This page displays a list of uploaded log files from clients, featuring information such as the file name, date, associated user and/or realm, client access component type, and the log node.



The system does not preserve uploaded logs when you upgrade the system software. To preserve the logs, you may archive them using options in the Maintenance > Archiving > Archiving Servers page of the admin console. You can also set the log-related SNMP traps to capture log events during the log upload using options in the System > Log/Monitoring > SNMP page of the admin console.

To view client log upload details:

1. In the admin console, choose **System > Log/Monitoring > Client Logs > Uploaded Logs management** page.

[Uploaded Log Listing Page](#) shows the management page.

2. (Optional) Refresh uploaded client log details by clicking the Refresh Logs button.
3. (Optional) View or save an uploaded log by clicking on its respective link.
4. (Optional) Delete an uploaded log by clicking the trash can icon in the right side of the log's column. Note that once you delete a log from a node, those logs are lost.

Uploaded Log Listing Page

Configuring SNMP

If you prefer, you can use a third-party SNMP manager, such as HP OpenView, to monitor system health. The system supports SNMP v2c and SNMPv3.

From 9.1R9 release, the system supports two users to be registered with an SNMP engine with different authentication and privilege settings.

To configure the SNMP agent:

1. Select **System > Log/Monitoring**.
2. Click the **SNMP** tab to display the **SNMP configuration page**.

[SNMP Configuration Page - Ivanti Connect Secure](#) shows the configuration page for Ivanti Connect Secure.

3. Complete the configuration as described in [Table](#).
4. Save the configuration.

[SNMP Configuration Page - Ivanti Connect Secure](#)

Log/Monitoring > SNMP

SNMP

Events | User Access | Admin Access | Sensors | Client Logs | **SNMP** | Statistics

MIB File

You must download the [Pulse Secure MIB file](#) and install it in your SNMP manager application to monitor the device.

SNMP Version data

SNMP Version:
 v2c v3

Agent Properties

SNMP Queries:
 SNMP Traps:

System Name:
 System Location:
 System Contact:
 Community:

Trap Thresholds

Set thresholds for traps.

Check Frequency: seconds (60-1800 seconds)

Log Capacity: % Disk: %
 Users: % CPU: %
 Physical Memory: % Meeting Users: %
 Swap Memory (Virtual Memory): %

To monitor the device for memory starvation condition it is recommended to use 'Virtual Memory' traps as Physical memory traps may get generated even if the device is not showing symptoms of memory starvation.

Optional traps

Critical Log Events
 Major Log Events

[Save Changes](#)

SNMP Trap Servers

Specify the servers to which the device will send any traps it generates.

10 records per page Search:

Hostname/IP Address (IPv4/IPv6)	Port	Community (optional)	
<input type="text"/>	<input type="text"/>	<input type="text"/>	Add
<input type="text"/>	<input type="text"/>	<input type="text"/>	
<input type="text"/>	<input type="text"/>	<input type="text"/>	
<input type="text"/>	<input type="text"/>	<input type="text"/>	
<input type="text"/>	<input type="text"/>	<input type="text"/>	

← Previous 1 Next →

The following table lists the SNMP Configuration Settings:

Settings	Guidelines
MIB File	Use the Ivanti MIB file link to download the device management information base MIB file. You add this file to your SNMP manager configuration.
SNMP Version	Select your SNMP server version: v2c v3

Settings	Guidelines
Agent Properties	
SNMP Queries	Select to support SNMP queries. Selecting this option enables the SNMP Diagnostic Log utility in the Troubleshooting > Monitoring > Diagnostic Logs page.
SNMP Traps	Select to send SNMP traps. Selecting this option enables the SNMP Diagnostic Log utility in the Troubleshooting > Monitoring > Diagnostic Logs page.
System Name	Specify a system name.
System Location	Specify a location.
System Contact	Specify a system contact.
Community String	Required only for SNMPv2c. To query the system, your network management station must send it the community string. To stop the SNMP system, clear the community field.
SNMPv3 Configuration	
Username	Specify the SNMPv3 username. The User-Based Security Model (USM) is the default Security Module for SNMPv3. The system supports two users to be registered with an SNMP engine. Editing the SNMPv3 user attributes overwrite any already registered SNMPv3 user. The SNMPv3 user must have read-only access on all MIBs supported by the system. SNMPv3 user configuration attributes can also be used for SNMP traps.

Settings	Guidelines				
Security Level	Selection	Auth Protocol	Auth Password	Priv Protocol	Priv Password
	No Auth, NoPriv	-	-	-	-
	Auth, NoPriv	Select MD5 (HMAC-MD5-96) or SHA (HMAC-SHA-96).	Enter an authentication password. The password can contain any ASCII characters and must be at least 8 characters in length.	-	-
	Auth, Priv	Select MD5 (HMAC-MD5-96) or SHA (HMAC-SHA-96).	Enter an authentication password. The password can contain any ASCII characters and must be at least 8 characters in length.	Select either CBC-DES or CFB-AES-128.	Enter a privacy password. The password can contain any ASCII characters and must be at least 8 characters in length.
Trap Thresholds	Setting a threshold value to 0 disables that respective trap.				
Check Frequency	Specify the frequency in seconds for sending traps. The default is 180 seconds.				
Log Capacity	Specify the percent of log space used. The default is 90%.				
Users	Specify the percent of user capacity used. The default is 100%.				
Physical Memory	Specify the percent of physical memory used. The default is 0 (not reported).				

Settings	Guidelines
Swap Memory (Virtual Memory)	Specify the percent of swap memory used. The default is 0 (not reported). We recommend you monitor swap memory to alert you to potential memory issues. The threshold for traps for physical memory usage might be reached even if the system is not experiencing any difficulties.
Disk	Specify the percent of disk utilization. The default is 80%.
CPU	Specify the percent of CPU utilization. The default is 0 (not reported).
Optional Traps	
Critical Log Events	Send traps when the system logs critical events.
Major Log Events	Send traps when the system logs major events.
Save SNMP Settings?	Click Save Changes to update the SNMP agent configuration. The page is refreshed and displays the SNMP engine ID. If the configuration is changed to move from SNMP v2c to SNMP v3, the system generates and displays two engine IDs.
SNMP Servers	
Hostname / IP address	Specify the hostname or IP address for the SNMP servers to which the system will send any traps it generates.
Port	Specify the port for the SNMP server. Typically, SNMP uses port 162.
Community (v2c) / User (v3)	Specify the community/user string (if necessary).

Keep the following configuration tips in mind when you configure your SNMP manager to listen for this SNMP agent:

- Add the Ivanti MIB file to the SNMP manager configuration.
- If using SNMPv2c, the community string configuration for the SNMP manager and SNMP agent must match.
- If using SNMPv3, the SNMPv3 user configuration for the SNMP manager and the SNMP agent must match.

- If using SNMPv3, you must specify the Authoritative Engine ID for SNMPv3 traps that was generated when you saved the SNMP agent configuration.

The following table is a reference of MIB objects for the system. Some objects apply only to Ivanti Connect Secure.

The following table lists the MIB Objects:

Object	Description
pulsesecure-gateway	This file defines the private Ivanti MIB extensions.
logFullPercent	Returns the percentage of available file size filled by the current log as a parameter of the logNearlyFull trap.
signedInWebUsers	Returns the number of users signed in through a Web browser.
signedInMailUsers	Returns the number of users signed in through a mail.
blockedIP	Returns the IP address-blocked due to consecutive failed login attempts-sent by the iveTooManyFailedLoginAttempts trap. The system adds the blocked IP address to the blockedIPList table.
authServerName	Returns the name of an external authentication server sent by the externalAuthServerUnreachable trap.
productName	Returns the licensed product name.
productVersion	Returns the software version.
fileName	Returns the file name sent by the archiveFileTransferFailed trap.
iveCpuUtil	Returns the percentage of CPU used during the interval between two SNMP polls. This value is calculated by dividing the amount of CPU used by the amount of CPU available during the current and previous SNMP polls. If no previous poll is available, the calculation is based on the interval between the current poll and system boot.

Object	Description
iveMemoryUtil	Returns the percentage of memory utilized by the system at the time of an SNMP poll. The system calculates this value by dividing the number of used memory pages by the number of available memory pages.
iveConcurrentUsers	Returns the total number of users logged in.
clusterConcurrentUsers	Returns the total number of users logged in for the cluster.
iveTotalHits	Returns the total number of hits to the system since last reboot. It includes total values from iveFileHits, iveAppletHits, and iveWebHits.
iveFileHits	Returns the total number of file hits to the system since last reboot. Incremented by the Web server with each GET/POST corresponding to a file browser request.
iveWebHits	Returns the total number of hits by means of the Web interface since last reboot. Incremented by the Web server for each http request received by the system, excluding file hits, and applet hits.
iveAppletHits	Returns the total number of applet hits to the system since last reboot. Incremented by the Web server for each GET request for a Java applet.
ivetermHits	Returns the total number of terminal hits to the system since last reboot.
iveSAMHits	Returns the total number of SAM (Secure Application Manager) hits to the system since last reboot.
iveNCHits	Returns the total number of NC (Network Connect) hits to the system since last reboot.
logName	Returns the name of the log (admin/user/event) for the logNearlyFull and iveLogFull traps.

Object	Description
iveSwapUtil	Returns the percentage of swap memory pages used by the system at the time of an SNMP poll. The system calculates this value by dividing the number of swap memory pages used, by the number of available swap memory pages.
diskFullPercent	Returns the percentage of disk space used in the system for the iveDiskNearlyFull trap. The system calculates this value by dividing the number of used disk space blocks by the number of total disk space blocks.
blockedIPList	Returns a table with the 10 most recently blocked IP addresses. The blockedIP MIB adds blocked IP addresses to this table.
ipEntry	An entry in the blockedListIP table containing a blocked IP address and its index (see IPEntry).
IPEntry	The index (ipIndex) and IP address (ipValue) for an entry in the blockedIPList table.
ipIndex	Returns the index for the blockedIPList table.
ipValue	A blocked IP address entry in the blockedIPList table.
logID	Returns the unique ID of the log message sent by the logMessageTrap trap.
logType	Returns a string sent by the logMessageTrap trap stating whether a log message is major or critical.
logDescription	Returns a string sent by the logMessageTrap trap stating whether a log message is major or critical.
ocspResponderURL	Returns the name of an OCSP responder.
fanDescription	Returns the status of the system fans.
psDescription	Returns the status of the system power supplies.
raidDescription	Returns the status of the system RAID device.
iveTemperature	Returns the temperature of MAG application blade. Other platforms such as ICS and IPS will return 0.

Object	Description
iveVPNTunnels	Returns the number of concurrent Pulse IPsec and NC users.
iveSSLConnections	Returns the total number of SSL connections.
esapVersion	Active ESAP version.
vipChangeReason	Reason for the VIP node change.
processName	Process name.
iveTotalSignedInUsers	Returns the total number of users logged in for the cluster.
vpnACLSPercentage	Returns the percentage of system ACL entries reached.
vpnACLSCount	Returns the number of system ACL entries reached.
blockedIPv6	The IPv6 address that is blocked due to consecutive failed login attempts.
iveNamedUsers	The total number of Named User Licenses used for the IVE node.
namedUserStorePercent	The storage space occupied in the Named Users store.
iveLogNearlyFull	<p>The log file (system, user access, or administrator access) specified by the logName parameter is nearly full. When this trap is sent, the logFullPercent (%of log file full) parameter is also sent. You can configure this trap to be sent at any percentage. To disable this trap, set the Log Capacity trap threshold to 0%. The trap's default value is 90%.</p> <p>When SNMP traps are enabled, the iveLogNearlyFull and iveLogFull traps are sent when the log files are 90% full and 100% full respectively, even if the threshold is set to 0 (disabled).</p>

Object	Description
iveLogFull	<p>The log file (system, user access, or administrator access) specified by the logName parameter is completely full.</p> <p>When SNMP traps are enabled, the iveLogNearlyFull and iveLogFull traps are sent when the log files are 90% full and 100% full respectively, even if the threshold is set to 0 (disabled).</p>
iveMaxConcurrentUsersSignedIn	<p>Maximum number or allowed concurrent users are currently signed in. You can configure this trap to be sent at any percentage. To disable this trap, set the Users trap threshold to 0%. The trap's default value is 100%.</p>
	<p>Setting the iveMaxConcurrentUsersSignedIn trap threshold to 0% only disables the threshold for the trap. The system continues to send SNMP traps generated from some other process in the system.</p>
iveTooManyFailedLoginAttempts	<p>A user with a specific IP address has too many failed sign-in attempts. Triggered when a user fails to authenticate according to the settings for the Lockout options on the Security Options tab.</p>
	<p>When the system triggers this trap, the system also triggers the blockedIP (source IP of login attempts) parameter.</p>
externalAuthServerUnreachable	<p>An external authentication server is not responding to authentication requests.</p> <p>When the system sends this trap, it also sends the authServerName (name of unreachable server) parameter.</p>
iveStart	The system has just been turned on.
iveShutdown	The system has just been shut down.
iveReboot	The system has just been rebooted.
archiveServerUnreachable	The system is unable to reach the configured archive server.

Object	Description
archiveServerLoginFailed	The system is unable to log into the configured archive server.
archiveFileTransferFailed	The system is unable to successfully transfer files to the configured archive server. When the system sends this trap, it also sends the fileName parameter.
iveRestart	Supplies notification that the system has restarted according to the administrator's instruction.
iveDiskNearlyFull	Supplies notification that the system disk drive is nearly full. When the system sends this trap, it also sends the diskFullPercent parameter. You can configure this trap to be sent at any percentage. To disable this trap, set the Disk trap threshold to 0%. This trap's default value is 80%.
iveDiskFull	Supplies notification that the system disk drive is full.
logMessageTrap	The trap generated from a log message. When the system sends this trap, it also sends the logID, logType, and logDescription parameters.
memUtilNotify	Supplies notification that the system has met the configured threshold for memory utilization. To disable this trap, set the Physical Memory trap threshold to 0. The threshold is 0%, by default.
cpuUtilNotify	Supplies notification that the system has met the configured threshold for CPU utilization. To disable this trap, set the CPU trap threshold to 0. The threshold is 0%, by default.
swapUtilNotify	Supplies notification that the system has met the configured threshold for swap file memory utilization. To disable this trap, set the Swap Memory trap threshold to 0. The threshold is 0%, by default.
ocspResponderConnectionFailed	OCSP Responder cannot be connected.
iveFanNotify	Supplies notification that the status of the fans has changed.

Object	Description
ivePowerSupplyNotify	Supplies notification that the status of the power supplies has changed.
iveRaidNotify	Supplies notification that the status of the RAID device has changed.
iveClusterDisableNodeTrap (clusterName,nodeList)	Supplies the name of the cluster that contains disabled nodes, as well as a string containing the names of all disabled nodes. Node names are separated by white space in the string.
iveClusterChangedVIPTrap(vipType, currentVIP, newVIP)	Supplies the status of a virtual IP for the cluster. The vipType indicates whether the changed VIP was external or internal. The currentVIP contains the VIP prior to the change, and newVIP contains the VIP after the change.
iveNetExternalInterfaceDownTrap (nicEvent)	Supplies the type of event that brought down the external interface. The nicEvent parameter can contain values of "external" for an external event and "admin" for an administrative action.
iveClusterDeleteTrap(nodeName)	Supplies the name of the node on which the cluster delete event was initiated.
iveNetInternalInterfaceDownTrap (nicEvent)	Supplies the type of event that brought down the internal interface. The nicEvent parameter can contain values of "external" for an external event and "admin" for an administrative action.
iveNetManagementInterfaceDownTrap (nicEvent)	Supplies the type of event that brought down the management port. The nicEvent parameter can contain values of "external" for an external event and "admin" for an administrative action.
iveTemperatureNotify	IVE Temperature is above threshold.
iveVIPNodeChanged	Notifies that VIP node has changed. nodeName is the new node which is hosting the VIP. vipChangeReason specifies the reason for the change.
iveProcessesNearMaxLimit	The count of processes (by processName) is about to reach maximum limit.

Object	Description
iveProcessesReachedMaxLimit	The count of processes (by processName) has reached maximum limit.
iveACLsNearMaxLimit	The percentage of ACL entries has reached maximum supported limit.
iveACLsCrossedMaxLimit	The count of ACL entries has crossed maximum supported limit.
iveTooManyFailedLoginAttemptsIPv6	Too many failed login attempts from IPv6 address.
iveMaxNamedUsersSignedIn	Maximum number of named users signed in.
iveNamedUsersStoreNearlyFull	Named user storage reached the limit.

Configuring Syslog

If desired, you can configure the system to send logs to a syslog server.

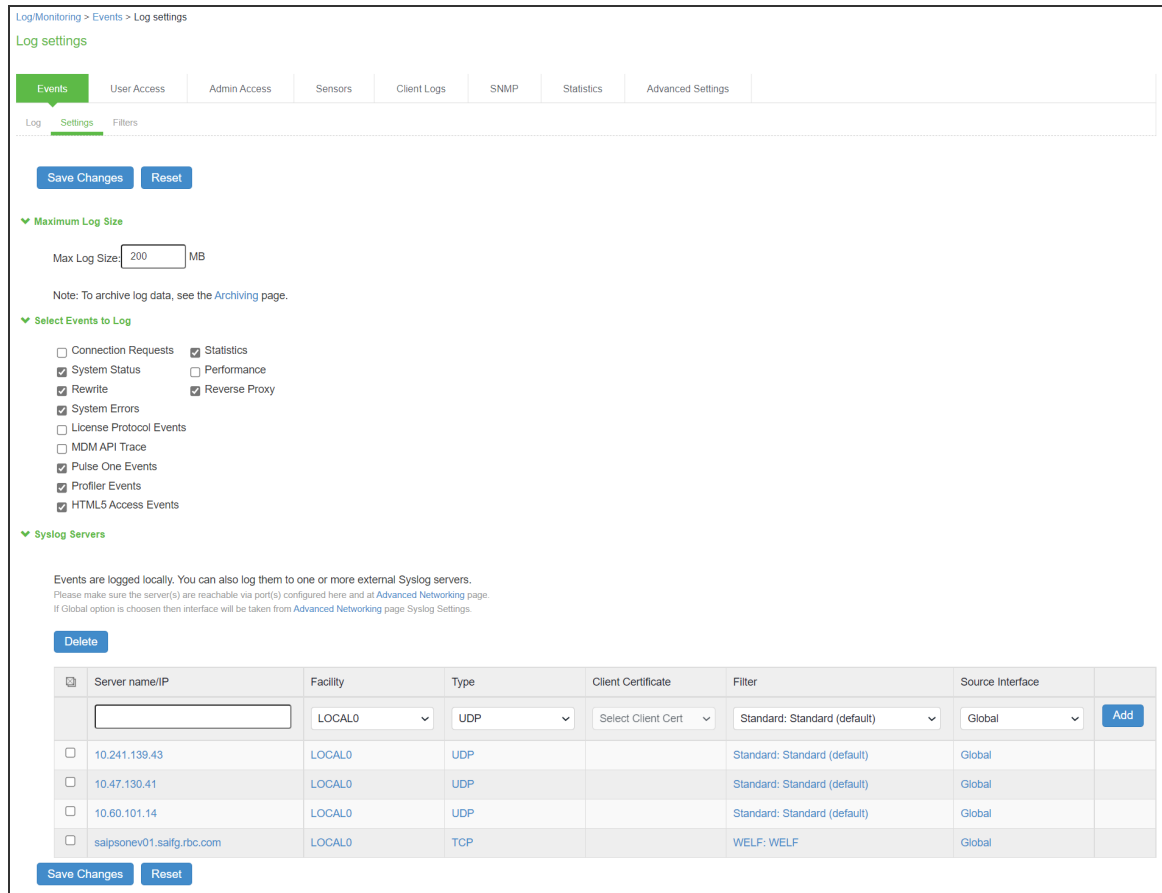
To configure reporting to a syslog server:

1. Select **System > Log/Monitoring**.
2. Click the **Settings** tab to display the configuration page. [Syslog Server Configuration Page - Ivanti Connect Secure](#) shows the configuration page for Ivanti Connect Secure. Specify the maximum log size and select the events to be logged. Specify the server configuration as described in [Table](#) and click **Add**. You can specify multiple syslog servers.
3. Save the configuration.



- To enable syslog reporting for each local log category, you must perform this procedure on each local log tab: Events, User Access, and Admin Access.
- ICS|IPS sends syslogs to remote syslog server (UDP|TCP|TLS) in compliance with Syslog RFC5424 (<https://tools.ietf.org/html/rfc5424>)

Syslog Server Configuration Page - Ivanti Connect Secure



The following table lists the Syslog Server Configuration Guidelines:

Settings	Guidelines
Server name/IP	Specify the fully qualified domain name or IP address for the syslog server.
	If you select TLS from the Type list, the server name must match the CN in the subjectDN in the certificate obtained from the server.
Facility	Select a syslog server facility level (LOCAL0-LOCAL7). Your syslog server must accept messages with the following settings: facility = LOG_USER and level = LOG_INFO.
Type	Select the connection type to the syslog server. You can select: UDP (User Datagram Protocol) - A simple non-secure transport model. TCP (Transmission Control Protocol) - A core protocol of the Internet Protocol suite (IP), but lacks strong security. TLS (Transport Layer Security) - Uses cryptographic protocols to provide a secure communication.

Settings	Guidelines
Client Certificate	(optional) If you select TLS from the Type menu and your remote syslog server requires client certificates, select the installed client certificate to use to authenticate to the syslog server. Client certificates are defined in the Configuration > Certificates > Client Auth Certificates page. Client certificates must be installed on the device before they can be used. There is no fallback if a connection type fails.
Filter	Select a filter format. Any custom filter format and the following predefined filter formats are available: Standard (default) -This log filter format logs the date, time, node, source IP address, user, realm, event ID, and message. WELF -This customized WebTrends Enhanced Log Format (WELF) filter combines the standard WELF format with information about the system realms, roles, and messages. WELF-SRC-2.0-Access Report -This filter adds access queries to the customized WELF filter. You can use this filter with NetIQ's SRC to generate reports on user access methods.
Source Interface	Select the source port type for the syslog server: <ul style="list-style-type: none"> • Global • External • Internal • Management <p>Ensure the servers are reachable through port configured in the Advanced Networking page on the Admin UI.</p>

Configuring Advanced Settings

This option helps to configure fault tolerance on each configured TCP and TLS syslog server available. Fault tolerance is supported only for TCP and TLS syslog servers. UDP syslog servers cannot be configured for fault-tolerance. This functionality helps the syslog server to recover the logs lost during a disconnect. The administrator can configure fault-tolerance on syslog servers by enabling this option from the admin UI. ICS/IPS reads the lost pending logs during a disconnect from the log disk and transports them to the syslog server on a reconnect. Fault tolerance is supported only for the syslog servers configured under the following log-types:

- Events

- User Access
- Admin Access



Fault tolerance is node-specific. In case of clusters, the setting needs to be enabled/disabled by logging into each of the cluster members. Fault tolerance is supported only for TCP and TLS syslog servers. UDP syslog servers cannot be configured for fault tolerance.

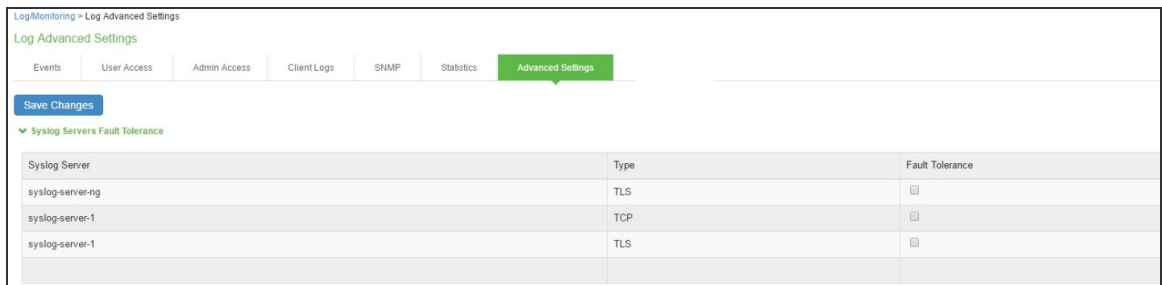
To configure advance settings to a TCP and TLS syslog server:

1. Select **System > Log/Monitoring**.
2. Click the **Advance Settings** tab to display the configuration page.
[Log Events Settings Configuration Page](#) shows the configuration page.
3. Complete the configuration as described in [Table](#).
4. Save the configuration.



This feature is limited to configuring fault tolerance settings of an existing syslog server; and cannot be used to create or delete a new syslog server.

Log Events Settings Configuration Page



The following table lists the Advanced Settings:

Settings	Guidelines
Syslog Server Fault Tolerance	
Syslog Server	Lists the existing Syslog servers.
Type	Specifies if the Syslog server is a TLS or TCP type.


Settings	Guidelines
Fault Tolerance	Tolerates the loss of network connection to a TCP/TLS syslog server for a brief period (maximum of 4 hours) by sending the logs missed during the disconnect time. Click the checkbox to enable this option. Fault-tolerance is disabled by default on any syslog server.

Displaying System Status

The System Status page is a dashboard of system version information, system capacity utilization, uptime, and summary user information. The System Status page is the "home" page that is displayed when you log into the admin console as an administrator. To navigate to the System Status page from other admin console pages, select **System > Status**.

The following figure shows the configuration page for Ivanti Connect Secure. The table that follows describes the numbered figure callouts.

System Status Page - Ivanti Connect Secure



Callout	Description
---------	-------------

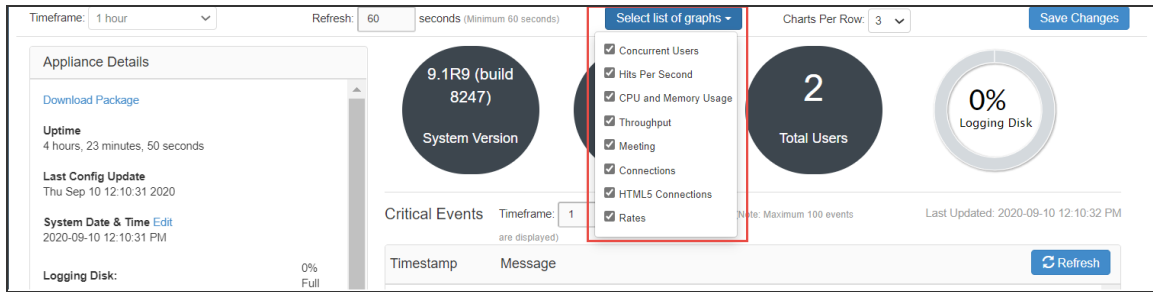
1	Click the Critical Events link to display a new window with a table of the last 10 critical system events.
2	Click the Page Settings link to display a new window with the System Status Settings page shown in System Status Settings Configuration Page
3	Click the System Version Download Package link to download the software version running on the system. You might do this when you need to synchronize software on another node to the software version running on this system.
4	Click the System Date and Time Edit link to display the System Date and Time configuration page. See Configuring the System Date and Time.
5	Click a System Capacity Utilization report Edit link to display a new window with controls to customize the appearance of the report graphs.
6	Click a System Capacity Utilization report Download link to download graph data in XML format.
7	Click an Enforcer Status link to navigate to its configuration page.

The following table lists the Licenses and Total Users - Ivanti Connect Secure:

Item	Description
Max Licensed Users	Displays the maximum number of licensed users by supported platform type.
Signed-In Users	Displays the number of signed-in users.
Signed-In Mail Users	Displays the number of signed-in mail users.
Concurrent Connections for Authorization only Access	Displays the concurrent connections for authorization only access.
ActiveSync Connections	Displays the number of ActiveSync connections.

The following figure shows the System Status Settings configuration page. The settings configuration page for Ivanti Connect Secure is similar.

System Status Settings Configuration Page



You can use this page to select the reports displayed on the System Status page, as well as data properties, such as the time dimension and refresh rate.

The following reports are available:

- **Concurrent Users** - Shows a count of users signed into the system. In clustered environments, the graph includes lines that display:
 - the number of local users signed into the node selected from the list
 - the number of concurrent users signed into the entire cluster.
 - In 9.1R8, added L4 access type (PSAM) and Clientless access type (Browser) logins as non-tunnel users.
- **Hits per Second** - Shows a count of hits currently being processed by the system. In a clustered environment, you may select a node from the list to determine which node's data is displayed in the graph. The graph includes three lines: total number of hits, number of Web hits, and number of client/server hits.
- **CPU and Memory Usage** - Shows the percentage of the CPU and memory being used. In a clustered environment, you may select a node from the list to determine which node's data is displayed in the graph.
- **Throughput** - Shows the amount of data (in KB) being processed. In a clustered environment, you may select a node from the list to determine which node's data is displayed in the graph. The graph includes four lines: external in, external out, internal in, and internal out.
- **Connections** - Shows a count of concurrent SSL connections.
- **HTML5 Connections** — Shows the traffic on the HTML5 RDP, HTML5 SSH, and HTML5 Telnet connections for Basic and Advanced solution types.
- **Rates** - Shows the rate of attempted logins, successful logins, and Host Checker updates.

Displaying Hardware Status

You can use the Maintenance > System > Platform page to display the hardware health status, including information about hard drives, fans, and power supplies.

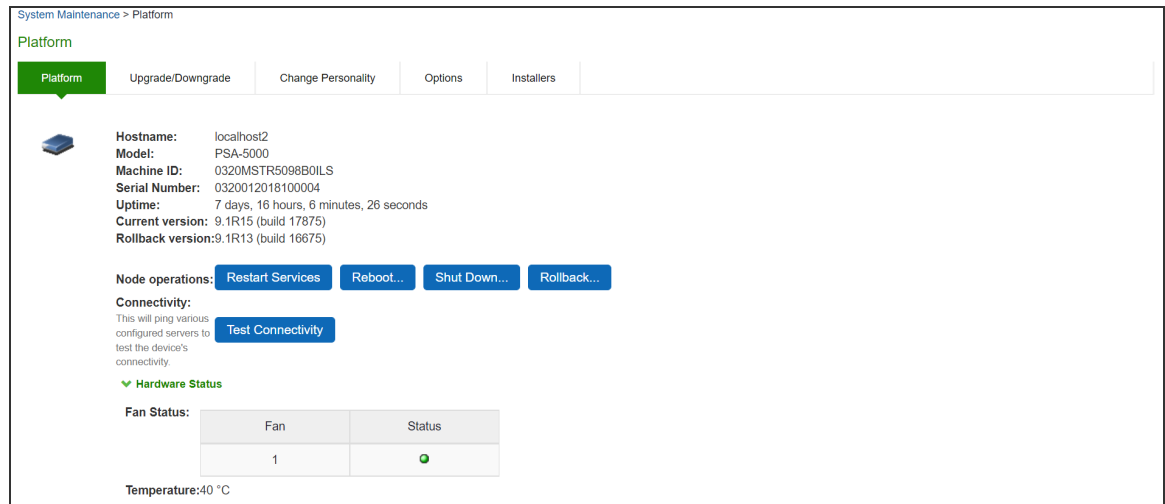
To display hardware health status:

1. Select **Maintenance > System Maintenance > Platform** to display the System Maintenance page.

[System Maintenance Page - Ivanti Connect Secure](#) shows the system maintenance page for Ivanti Connect Secure.

2. Review the hardware status information described in [Table](#).

System Maintenance Page - Ivanti Connect Secure



The following table lists the Hardware Status Information:

Hardware Component	Status Message
Hard Disk Status	Displays a health statement for the device disk drive. See Table 156 and Table 157 for details.
Fan Status	Displays a health statement for the device fan(s).
Power Supply	Displays a health statement for the device power supply.

The following table lists the RAID status and hard drive status. Depending on your system, you may or may not see all these possible statuses.

RAID Status	Drive 1	Drive 2
Hard Disk RAID is operational	Active	Active
Hard Disk RAID is in single drive mode	Missing	Active
Hard Disk RAID is in single drive mode	Active	Missing
Hard Disk RAID has failed	Failed	Active
Hard Disk RAID has failed	Active	Failed
Hard Disk RAID is in the process of recovering	Active	Reconstructing
Hard Disk RAID is in the process of recovering	Reconstructing	Active
Hard Disk RAID is in the process of recovering	Active	Verifying
Hard Disk RAID is in the process of recovering	Verifying	Active
Hard Disk RAID status is unknown	Unknown	Active
Hard Disk RAID status is unknown	Active	Unknown
Hard Disk RAID status is unknown	Unknown	Unknown
Not available	n/a	n/a

Using Software RAID PSA7000

This section describes the use of software RAID on the PSA7000.

All hard disks are encrypted with AES128 using a random generated key.

Overview of Software RAID on the PSA7000

PSA-7000 hardwares have two hard disks but, have no hardware RAID controller. RAID functionality is enabled through the software RAID available in Linux. Software is used to create RAID devices in the RAID 1 layer. These individual disk partitions can also be used as regular partitions. All partitions, which include boot, root, data, var, tmp, swap, and so on are created out of the software RAID. Software RAID does not affect any of the other hardware models except PSA7000. It works the same way either with hardware RAID or no RAID at all.

Configuring RAID Controller on the PSA7000

To configure a RAID controller on the PSA7000, disks are hot plugged on the PSA7000. Configuring RAID on the PSA7000 also involves some manual configuration steps in the admin console. RAID configuration is carried out by the following steps:

To remove a disk from the second slot in the system

1. Select Menu option: **4. System Operations**
2. Select Menu option: **20. Manage RAID**
3. A prompt appears: **Are you sure you want to manage RAID? (y/n)**
4. Enter y for the prompt.
5. Choose Menu option: **2. Remove Disk2**. This will detach the disk in the second slot from software RAID.
6. Physically unplug the disk from the second slot.

To insert a disk to the second slot in the system:

1. Physically plug in the disk in the second slot.
2. Select Menu option: **4. System Operations**
3. Select Menu option: **20. Manage RAID**
4. A prompt appears: **Are you sure you want to manage RAID? (y/n)**
5. Enter y for the prompt.
6. Choose Menu option: **4. Add Disk2**. This will attach the disk in the second slot to the software RAID.

Checking RAID Statuses

To check the status of RAID and individual disks:

1. Go to **system > platform** page of the web interface.

Status of RAID and individual disks are displayed. The following table shows the hard disk status and the hard disk RAID status.

The following table lists the Hard Disk and Hard Disk RAID Statuses:

Status	Guidelines	
Hard Disk	Active	A disk that is present and part of RAID
Inactive	A disk that is present but not part of RAID	
Missing	A disk that is removed.	
Hard Disk RAID	Operation	Both disks are active and part of the software RAID
Recovering	Both disks are active and syncing with each other	
Failed	Both disks are active but one of them is not part of RAID	

In the single drive mode, only one disk is active. The other disk is either missing or inactive.

LCD Display

This section describes the addition of LCD to ICS devices.


Overview of adding LCD for ICS

The addition of an LCD screen allows field technicians to quickly gauge the health of the system without logging into the device. The buttons on the LCD panel allow navigation through the display menus. The directional buttons are used to access the menu modes and find device information. The LCD can display two line of text. [LCD with Navigation Buttons](#) shows the LCD screen with navigation buttons.



LCD display is available for the ICS-7000 platform model only.

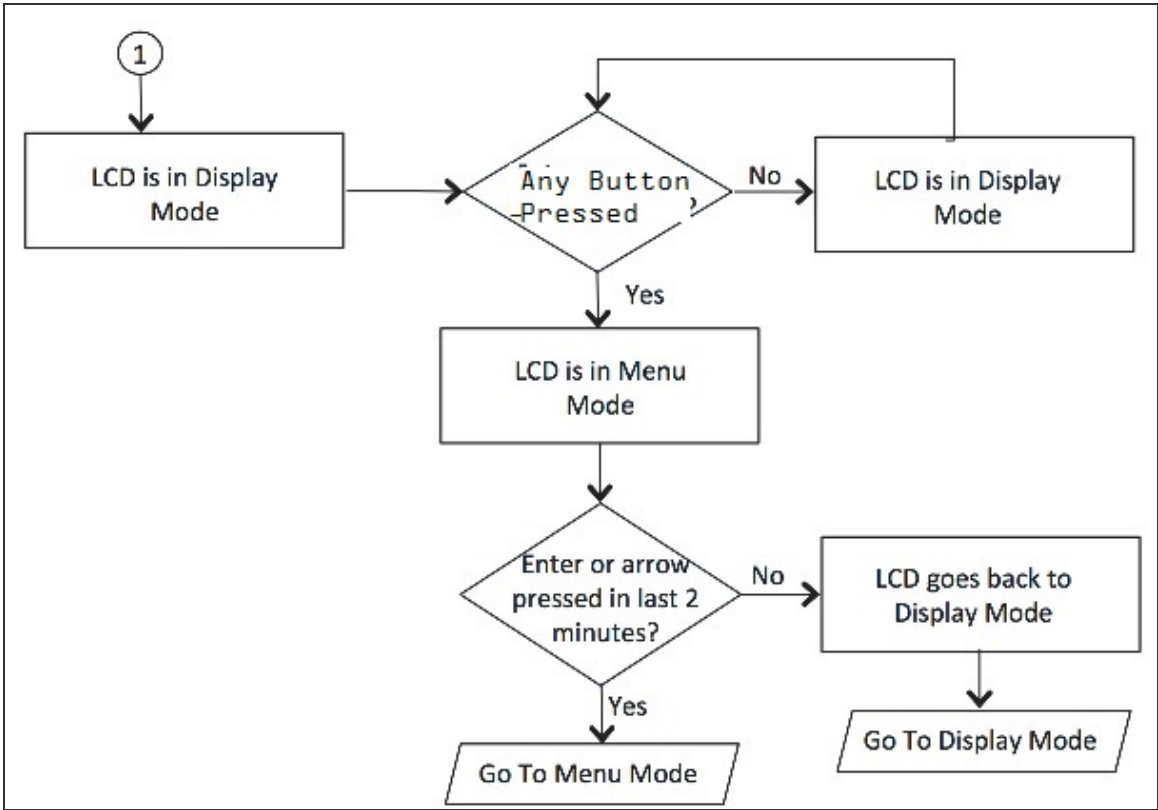
LCD with Navigation Buttons



Modes Supported by the LCD

The LCD supports two modes namely the display mode (default) and the menu mode. Pressing any button in the display mode will change the mode to menu mode. If a user presses the cancel button, the LCD immediately changes back to display mode and shows the appropriate state. The LCD remains in display mode. If the LCD is in menu mode and the user does not press any button for more than two minutes, then the LCD changes back to display mode. [Two Modes Supported by the LCD](#) shows the two modes supported by the LCD


Two Modes Supported by the LCD



Display Mode

The display mode describes the current state of the system, such as normal state or error conditions (e.g., fan speed and overheat). It represents the default status. The LCD goes into display mode after boot-up is complete. In display mode, the LCD is either set to NORMAL or shows a label that describes an error condition. If all systems are functioning normally, then the LCD shows NORMAL. The second line in the NORMAL state is used to show whether the appliance is configured as part of a cluster. The valid states in the display mode are Clustering OFF and Clustering ON. [Valid States in Display Mode](#) shows the two valid states in display mode.

Valid States in Display Mode



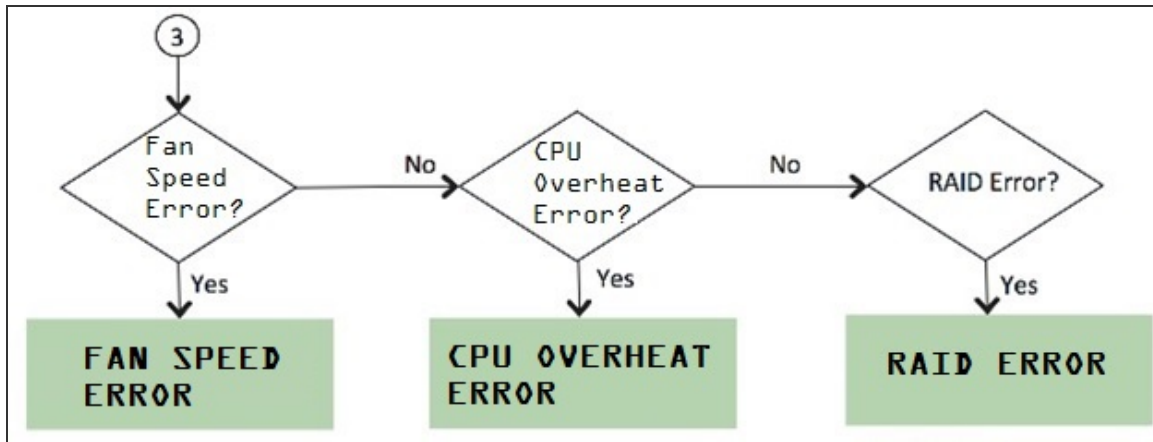
Detecting Error Conditions in Display Mode

If more than one error condition is detected, all error conditions will be displayed in sequence with a 2 second pause before switching to the next one. All error conditions need to be cleared before the status returns back to the NORMAL state. Error conditions include:

- Overheat
- Fan Failure
- RAID Errors

[Error Conditions](#) shows the various error conditions that are detected.

Error Conditions



To detect the error conditions in the display mode:

If there are any error conditions, they are automatically shown on the LCD screen when it is in the display mode. The types of errors displayed are: Fan Failures, CPU overheating and RAID errors. If there are multiple errors, they would be displayed in the order shown in [Error Conditions](#) with a two second pause between successive displays.

xZZZcx cmcc nxncxncxxxxkxkkv;dkdkkm The error message is automatically cleared when the underlying error condition is resolved. For example: the CPU overheat message disappears when CPU temperature is lowered. The user can enter the Menu mode at any point, even if an error message is being displayed.

Menu Mode

The menu mode is activated when the user presses any button. A single press of the button changes to menu mode and loads the last selected menu selection.

To view information in the menu mode:


1. Press any button. This puts the LCD into menu mode.
2. Press the right and left arrows keys to obtain the available system configuration data.
3. View information starting with the Internal IP and moving in a clockwise direction.
4. The menu screens loop back in a cycle.
5. Press Cancel at any point to exit to display mode.



Any button, even cancel will put the user in the menu mode.

[System Configuration Data Available in Menu Mode](#) show the available system configuration data available in menu mode.

System Configuration Data Available in Menu Mode



Displaying Active Users

You can use the Active Users page to display the system active users table and to perform administrative actions pertaining to active sessions.

The system active users table displays all users who have an active session (in contrast to the user's tables that appear on the authentication server configuration pages, which display session records for active and inactive sessions that were authenticated by the particular authentication server).

If a user signs in and is placed in a VLAN without an IP address, the table does not display an IP address under Signed in IP.

If there is a NAT device between the user's computer and the Infranet Enforcer, the table displays both the NAT device's IP address and the endpoint's virtual source IP address under Signed in IP. For example, if the NAT device's IP address is 10.64.9.26, and the endpoint's virtual source IP address is 192.168.80.128, the following information is displayed under Signed in IP: **10.64.9.26 (192.168.80.128 behind NAT)**.

To display the system Active Users page:

1. Select **System > Status**.
2. Click the **Active Users** tab to display the system active users page.
3. Use the controls described in the following table to perform administrative actions pertaining to active sessions.

The following table lists the Active Users Page:

Buttons	Administrative Actions
Update	<p>Refresh records displayed on the page: To refresh the page, click Update.</p> <p>To display a specific user, enter the username in the Show Users Named box and click Update. If you do not know the exact username, use the asterisks (*) as a wildcard character.</p> <p>To change the table size, enter a number in the Show N users' box and click Update.</p> <p>To sort the table of currently signed-in users and administrators, click a column header.</p>
Delete Session	<p>Select the check box next to the appropriate names and then click Delete Session to immediately delete the session. The user is signed out by your action.</p>
Delete All Sessions	<p>Use this option to immediately delete all sessions. Users are signed out by your action.</p> <p>If you want to sign out administrators, you must choose them individually and use the Delete Session button.</p>
Refresh Roles	<p>Manually evaluate all authentication policies, role-mapping rules, role restrictions, user roles, and resource policies for all currently signed-in users. Use this button if you make changes to an authentication policy, role-mapping rules, role restrictions, or resource policies and you want to immediately refresh the roles of all users.</p>

Displaying System Logs

This topic describes how to display local system logs.

Displaying Events Logs

The Events logs include system events, such as session timeouts, system errors and warnings, requests to check server connectivity, and system restart notifications. The local log viewer displays the most recent 5000 log messages (the display limit).

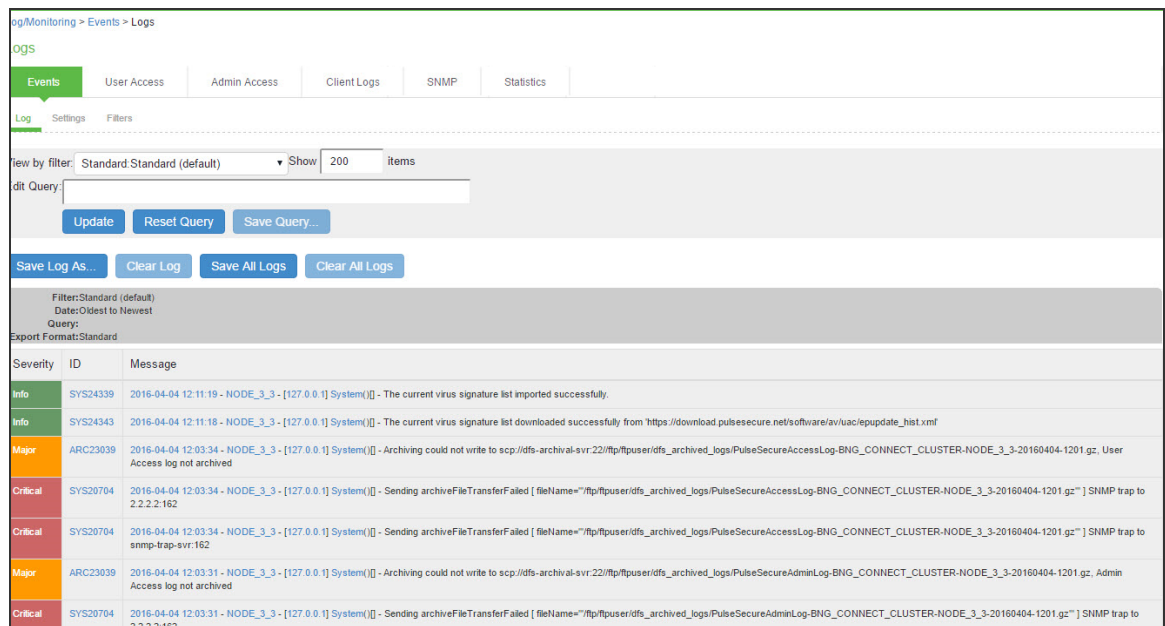
To display Events logs:

1. Sselect **System Log/Monitoring**.
2. Click the **Events** tab.
3. Click the **Log** tab to display the log page.

[Events Logs Page - Ivanti Connect Secure](#) shows the log page for Ivanti Connect Secure.

4. Use the features described in [Table](#) to examine log records or manage the log collection.

Events Logs Page - Ivanti Connect Secure



The following table lists the Log Management Features:

Controls	Description
Filter	<p>Select a filter format. Any custom filter formats and the following predefined filter formats are available:</p> <p>Standard (default)-This log filter format logs the date, time, node, source IP address, user, realm, event ID, and message.</p> <p>WELF-This customized WebTrends Enhanced Log Format (WELF) filter combines the standard WELF format with information about the system realms, roles, and messages.</p> <p>WELF-SRC-2.0-Access Report-This filter adds access queries to the customized WELF filter. You can use this filter with NetIQ's SRC to generate reports on user access methods.</p> <p>Format filters change only the data displayed (or columns exported), and do not affect the log data that has been collected.</p>
Query	<p>In the log display, several fields are hyperlinks. The hyperlinks function as dynamic queries on the local log collection. For example, if you click the log ID, the date, or an IP address or username, the log viewer queries the log collection for records that match the value you clicked, and redisplay the log collection. You can apply additional query filters by clicking additional hyperlinked values, essentially creating a Boolean AND query (for example, date AND IP address). Use the Reset Query button to clear the query filters and redisplay the unfiltered log collection.</p> <p>Use the Save Query button to save the dynamic log query as a custom filter. When you click the Save Query button, the system displays the Filters tab displays with the Query field prepopulated with the variables you selected from the log.</p> <p>Query filters change only the display (or rows exported), and do not affect the log data that has been collected.</p>

Controls	Description
Save Log As	<p>Save the local log collection to a file. We recommend you retain the system generated log name, which follows a consistent convention: juniper.logtype.nodename.log.</p> <p>The local log viewer displays the most recent 5000 log messages (the display limit). If the current log file contains fewer than 5000 log messages, older log messages from the backup log file are displayed, up to a total of 5000 log messages. This makes the log files appear as one, even though they are stored separately.</p> <p>When you save the log messages or use the FTP archive function, the backup log file is appended to the current log file, and is then downloaded as one log file. If the log files are not archived or saved by the time they are rolled over again, the oldest log messages (saved in the backup log file) are lost.</p>
Clear Log	<p>Clear the local log and log.old file.</p> <p>When you clear the local log, events recorded by the syslog server are not affected. Subsequent events are recorded in a new local log file.</p>
Save All Logs	<p>The Save All Logs button appears on the Events, User Access, and Admin Access tabs. When you click Save All Logs, the system generates a file that includes event, user access, admin access, and XML data for all of the system statistics and graphs shown on the Status > Overview page. After you click Save All Logs, you are prompted to download a file named pulsesecurelogs-graphs.tar.gz to your local host.</p>
Clear All Logs	<p>The Clear All Logs button appears on the Events, User Access, and Admin Access tabs. It clears event, user access, admin access, and XML data for all of the system statistics and graphs shown on the Status > Overview page. When you clear the local log, events recorded by the syslog server are not affected. Subsequent events are recorded in a new local log file.</p>

Displaying User Access Logs

The User Access logs include information about user access, such as the number of simultaneous users at each one-hour interval (logged on the hour) and user sign-ins and sign-outs. The local log viewer displays the most recent 5000 log messages (the display limit).

To display User Access logs:

1. Select **System > Log/Monitoring**.

2. Click the **User Access** tab.
3. Click the **Log** tab.
4. Use the features described in [Table](#) to examine log records or manage the log collection.

Displaying Admin Access Logs

The Admin Access logs include information about administrator actions, such as administrator changes to user, system, and network settings. It includes a log entry whenever an administrator signs in, signs out, or changes licenses on the appliance. The local log viewer displays the most recent 5000 log messages (the display limit).

To display Admin Access logs:

1. Select **System Log/Monitoring**.
2. Click the **Admin Access** tab.
3. Click the **Log** tab.
4. Use the features described in [Table](#) to examine log records or manage the log collection.

Using Log Filters

This topic describes how to use log filters.

Reviewing the Configuration of Predefined Log Format Filters

To view the configuration of predefined log format filters:

1. Select **System > Log/Monitoring**.
2. Click the **Events** tab.
3. Click the **Filter** tab to display the log filters page.

[Log Filters Page - Ivanti Connect Secure](#) shows the log filters page for Ivanti Connect Secure.

4. Click the hyperlinked name of the filter to display its configuration page. You cannot edit the predefined filter named Standard, but you may edit the predefined WELF filters and any other custom filters that appear in the list.

Log Filters Page - Ivanti Connect Secure

Log/Monitoring > Events > Log filters

Log filters

Events | User Access | Admin Access | Client Logs | SNMP | Statistics

Log | Settings | Filters

New Filter... Delete...

10 records per page

Search:

Filter	Query	Export Format
Standard (default)	Date: Oldest to Newest Query:	Standard
<input type="checkbox"/> WELF	Date: Oldest to Newest Query:	WELF
<input type="checkbox"/> WELF-SRC-2.0-Access-Report	Date: Oldest to Newest Query: id = ('VEB20174' or 'EML20825' or 'FBR20512' or 'FBR20503' or 'FBR20501' or 'FBR20536' or 'FBR20540' or 'JAV20023' or 'NWC23464' or 'NWC23465' or 'MTG20742' or 'MTG20749' or 'MTG20866' or 'MTG20869' or 'MTG20875' or 'MTG20877' or 'STA22721')	WELF

Previous 1 Next

Creating a Custom Log Collection Filter

If desired, you can create custom log collection filters to change the records displayed or exported. For example, it is common to see administrators use a filter for RADIUS accounting logs. This filter allows only the accounting log message, and it puts the entire message in a comma separated list. The order of the filtered message is: Date, Time, User, Realm, "List of Roles", NAS-ID, Acct-Status, Auth-Type, Attr-Value1, Attr-Value2, Attr-Value3.

Accounting attribute messages are different from authentication attribute messages in that the attribute name is not printed in the log message, but a comma is inserted for every attribute to be logged, even if it is not present.

To create a custom log collection filter:

1. Select **System > Log/Monitoring**.
2. Click the **Events** tab.
3. Click the **Filter** tab.
4. Click **New Filter** to display the configuration page. [New Filter Page - Ivanti Connect Secure](#) shows the configuration page for Ivanti Connect Secure.
5. Complete the configuration as described in [Table](#).
6. Save the configuration.

New Filter Page - Ivanti Connect Secure

The screenshot shows the 'New Log filter' configuration page. At the top, there are navigation tabs: 'Events' (selected), 'User Access', 'Admin Access', 'Sensors', 'Client Logs', 'SNMP', and 'Statistics'. Below these are sub-tabs: 'Log', 'Settings', and 'Filters' (selected). The main section is titled 'Filter' and contains the following elements:

- Filter Name:** A text input field.
- Make default for syslog and archiving filter selection
- Query:** A section with 'Start Date' (radio buttons for 'Earliest Date' and 'Latest Date (moving)', with 'Earliest Date' selected) and 'End Date' (radio buttons for 'Latest Date (moving)' and another option, with 'Latest Date (moving)' selected). Below these are date selection dropdowns.
- Query:** A large text area for entering the filter query.
- Export Format:** A section with 'Format' options: 'Standard' (selected), 'WELF', and 'Custom'. Below is a text area containing a sample format string:


```
%date% %time% - %node% - [%sourceip%] %user%(%realm%) [%role%] - %msg%
```
- Filter Variables Dictionary:** A panel on the right showing a list of variables: result, port, method, srcport, uri, and sbvtes. Each variable has a right-pointing arrow. At the bottom of this panel is a button labeled '< Insert Expression'.

At the bottom of the page are 'Save' and 'Cancel' buttons.

The following table lists the Filter Settings:

Settings	Guidelines
Filter Name	Specify a name that is helpful to you and other administrators in understanding usage for your custom filter.
Make default	Make the filter the default on syslog and archiving configuration pages.
Query	
Start Date	Enter a start date. Click Earliest Date to write all logs from the first available date stored in the log file.

Settings	Guidelines
End Date	Enter an end date. Click Latest Date to write all logs up to the last available date stored in the log file.
Query	Use the Filter Variables Dictionary to insert query expressions in the Query box. Enclose the query value in single quotes. For example, insert the query expression sourceip=. Then complete the expression by adding the value ' 192.168.0.1 '.
Export Format	Select an export format: Standard (default) -This log filter format logs the date, time, node, source IP address, user, realm, and message. WELF -This customized WebTrends Enhanced Log Format (WELF) filter combines the standard WELF format with information about the system realms, roles, and messages. Custom -Use the Standard as a template for your custom selection of columns to be included in exports (when log collections are saved to files).



Log query filters change only the data displayed (or rows exported). Log format filters change only the data displayed (or columns exported). Use of filters does not affect the log data that has been collected.

Example: Using the Source IP Address Filter

When drilling into logs to verify behavior or troubleshoot an issue with a dual-stack device, it is helpful to redisplay the log collection filtered on the IP address.

To filter on an IP address:

1. Select **System > Log/Monitoring**.
2. **Create the filter:**
 1. Select **User Access** and then **Filter**.
 2. Define the filter expression, name the filter, and click **Save**. In this example, we create a filter based on source IP address and name it IPv6_Address_Filter:Standard.
3. **Use the filter:**
 1. Select **Logs** to display the user logs table.

2. Under **View** by filter, select IPv6_Address_Filter:Standard, as shown in [Using IP Address Filters](#).
3. If desired, under Edit Query, edit the value of the sourceip= variable expression to filter on different source IP addresses.
4. Click **Update** to apply the filter and redisplay the log collection.

Using IP Address Filters

Displaying User Access Statistics

Every hour, the system logs the peak count of Web users in the previous hour. It displays the hourly counts for the past week on the Statistics page. It writes the report to the system log once a week.

To display user statistics:

1. In the admin console, select **System > Log/Monitoring**.
2. Click the **Statistics** tab to display the page.

[User Statistics Page - Ivanti Connect Secure](#) shows the configuration page for Ivanti Connect Secure.

3. Scroll the page to view the data.

[User Statistics Page - Ivanti Connect Secure](#)

Log/Monitoring > Monitoring Statistics

Monitoring Statistics

Events | User Access | Admin Access | Sensors | Client Logs | SNMP | **Statistics**

Signed-In Users
Hourly peak load of users

	Sunday 4/3/2016	Monday 4/4/2016	Tuesday 3/29/2016	Wednesday 3/30/2016	Thursday 3/31/2016	Friday 4/1/2016	Saturday 4/2/2016
12:00 am	5	3	0	0	0	7	8
01:00 am	5	2	0	0	0	7	7
02:00 am	4	1	0	0	0	5	7
03:00 am	4	1	0	0	0	5	7
04:00 am	4	1	0	0	0	5	5
05:00 am	4	1	0	0	0	5	5
06:00 am	3	1	0	0	0	4	5
07:00 am	4	3	0	0	0	3	3
08:00 am	5	5	0	0	0	4	2
09:00 am	5	8	0	0	0	5	3
10:00 am	7	9	0	0	1	7	4
11:00 am	8	8	0	0	6	8	4
12:00 pm	7	8	0	0	9	8	4
01:00 pm	9	0	0	0	9	8	4
02:00 pm	10	0	0	0	8	8	4

- Upgrading software clears all statistics. If you configure the system to log statistics hourly, however, older statistics are still available in the log file after an upgrade.

Troubleshooting Tools

Using the Admin Console Troubleshooting Tools

You can use the admin console troubleshooting tools to investigate user access issues and system issues. The following tools are available through the Maintenance > Troubleshooting pages:

- **Policy tracing** - Diagnose user access issues.
- **Simulation** - Ivanti Connect Secure only. Diagnose user access issues.
- **Session recording** - Ivanti Connect Secure only. Work with Global Support Center to diagnose user access issues.
- **Debug logs** - Work with Global Support Center to diagnose system issues.
- **tcpdump** - Sniff packet headers to diagnose networking issues.
- **Network troubleshooting commands** - Use standard network commands, such as ping, traceroute, NSlookup, and other commands to diagnose networking issues.
- **Kerberos debugging** - Diagnose issues with Kerberos communication.
- **System snapshots** - Work with Global Support Center to reproduce and diagnose system issues.
- **Remote debugging** - Enable Global Support Center to access your system directly to help you diagnose system issues.

If the admin console is unavailable, you can use the serial port console to perform some troubleshooting operations, such as use ping and traceroute commands, view logs, create system snapshots, and perform configuration rollbacks and factory resets.

Using Policy Tracing


It is common to encounter a situation where the system denies a user access to the network or to resources, and the user logs a trouble ticket. You can use the policy tracing utility and log to determine whether the system is working as expected and properly restricting access, or whether the user configuration or policy configuration needs to be updated to enable access in the user's case.

To create a policy trace log:

1. Select **Maintenance > Troubleshooting > User Sessions > Policy Tracing** to display the configuration page.

[Policy Tracing Configuration Page](#) shows the policy tracing configuration page for Ivanti Connect Secure.

Policy Tracing Configuration Page



2. Complete the configuration as described in the following table

The following table lists the Policy Trace Configuration Guidelines:


Settings	Guidelines
Record Trace File	
User	Specify the username to trace.
Source IP	Specify the source IP address if you know it. If you are able to provide the source IP address, the policy trace log can include events that occur before the user ID is entered into the system.
Realm	Select the realm to trace.

Settings	Guidelines
Events to Log	
Pre-Authentication	Logs events related to evaluation of realm rules.
Authentication	Logs events related to authentication.
Role Mapping	Logs events related to role mapping.
Web Policies	Logs events related to web policies.
File Policies	Logs events related to file policies.
SAM Policies	Logs events related to SAM policies.
Terminal Services Policies	Logs events related to terminal services.
VPN Tunneling Policies	Logs events related to VPN tunneling.

3. Click **Start Recording**.

[Policy Tracing Page During Recording](#) shows the policy tracing page with the recording indicator.


Policy Tracing Page During Recording



4. Initiate the action you want to trace, such as a user sign in.
5. Click **View Log** to display the policy trace results log.
6. Click **Stop Recording** when you have enough information.

[Policy Tracing Results Page](#) shows the page with policy trace results.

Policy Tracing Results Page



The following table describes options for managing the policy trace results log file.

The following table lists the Post-Trace Options:

Control	Guidelines
Delete Trace	Under Events to Log, click Delete Trace to clear the results displayed on this page.
Update	Specify a number of rows to display and click Update to change the number of rows that are displayed.
Save Log As	Click this button to save the trace results log to a file. This is useful particularly when you are working with the Global Support Center to troubleshoot a case.
Clear Log	Click this button to clear the log file from the system.

Using the Simulation Utility

Ivanti Connect Secure allows you to troubleshoot problems by simulating the events causing the problem. Using the Maintenance > Troubleshooting > User Sessions> Simulation page, you can create virtual user sessions without requiring actual end users to sign in to the device and recreate their problems. In addition, you can also use the Simulation tab to test new authentication and authorization policies before using them in a production environment.

To use the simulator, you must specify which events you want to simulate (for example, you can create a virtual session in which "John Doe" signs into the "Users" realm at 6:00 AM from an Internet Explorer browser). Then, you must specify which events you want to record and log in the simulation. You can log three major types of events to the simulation log:

- **Pre-Authentication** - The system events that are captured will not include any other system related events. Events are merely used as a filtering mechanism to reduce the number of logs and highlight the problem.
- **Role Mapping** - The system events that are captured will not include any other system related events. Events are merely used as a filtering mechanism to reduce the number of logs and highlight the problem.
- **Resource Policies** - The system events that are captured will not include any other system related events. Events are merely used as a filtering mechanism to reduce the number of logs and highlight the problem.

To simulate a user session:

1. In the admin console, choose **Maintenance > Troubleshooting > User Sessions > Simulation**.

[Simulation Configuration Page](#) shows the configuration page for Ivanti Connect Secure.

2. In the Query Name field, enter a name for the query.
3. In the Username field, enter the username of the user whose experience you want to simulate. Note that you may use a wildcard character (*) in place of a username.
4. From the Realm drop-down menu, select the realm of the user whose experience you want to simulate.
5. If you want to determine whether to apply a specific type of resource policy to a user's session, enter the specific resource you want to simulate in the Resource field and select a policy type from the Resource drop-down list. Then:

- If you want to determine whether a user can successfully sign in to the device, select the Pre-Authentication check box.
- If you want to determine whether a user can successfully map to a specific role, select the Role Mapping check box. Note that this option controls whether role mapping results are logged to the simulator log, not whether to run role mapping rules. The system always runs role mapping rules, even if you do not select this check box.
- Specify the types of policies you want to log using the check boxes in the Events to Log section.

For example, if you want to test whether a user can access the Yahoo web site, enter "http://www.yahoo.com" in the Resource field, select Web from the drop-down list, and select the Access check box in the Events to Log section.

6. In the Variables section, use a combination of text and variables to create a custom expression that reflects the exact same values as in the real session of the user who is facing a problem. For example, if you want to create a session in which the user signs in to the device at 6:00 AM, enter "time = 6:00 AM" in the Variables field. For complete instructions on how to create a custom expression. You may also view the syntax for a given variable by clicking the arrow next to it in the Variables Dictionary.

If you fail to create a custom expression that includes the virtual user's IP address, the system uses your current IP address instead. Also note that if you use the role variable to specify the role of the virtual user (for example, role="Users"), the system ignores results from role mapping rules and assigns the virtual user to the role(s) you specify.

7. Choose one of the following options:
 - **Run Simulation**-Runs the specified simulation and creates an on-screen log file.
 - **Save Query**-Saves the query.
 - **Save Query and Run Simulation**-Runs the specified simulation and also saves it for later use.
8. After running the simulation, choose **Save Log As** to save the simulation results to a text file.

Simulation Configuration Page

Troubleshooting

User Sessions | Monitoring | Tools | System Snapshot | Remote Debugging

Policy Tracing | **Simulation** | Session Recording | Virtual Desktop

View: New

Query name: Name this query for future use.

Username:

Realm: (Select a realm)

Resource: - Select a resource type -

Events To Log

- Pre-Authentication
- Role Mapping
- IF-MAP
- Web Policies**
 - Access
 - Rewriting
 - SAML
 - Cross Domain Access
 - Caching
 - Web Proxy
 - Launch JSAM
 - Compression
 - Compression (Windows)
 - Compression (UNIX/NFS)
 - Java
 - Protocol
 - Compression
 - Kerberos/NTLM/Basic Auth
 - SSO Post/SSO Headers
 - Encoding
 - Rewriting Filters
 - Persistent Cook
- File Policies**
 - Windows
 - Compression (Windows)
 - UNIX/NFS
 - Compression (UNIX/NFS)
 - Windows Credentials
- SAM Policies
- Telnet/SSH Policies
- Terminal Services Policies
- HTML5 Access Policies
- VPN Tunneling Policies

Variables

Variables: Only 1 variable/value pair per line.

Variables Dictionary

- certAttr: C
- certAttr.altName: directoryName
- certAttr.serialNumber
- certDNText
- certIssuerDNText
- groups
- hostCheckerPolicy
- loginHost
- loginTime
- loginURL
- networkif
- role

< Insert Expression

Save or Run Simulation?

Run Simulation | Save Query | Save Query and Run Simulation

Using the Session Recording Utility

You can use the Session Recording utility to record a trace file that lists a user's actions when accessing a resource or connecting to a client/server application. You do this to troubleshoot issues users might report regarding the Web access or client access.

When you start recording a trace file, the system signs out the specified user and then starts recording all user actions after the user signs in again and is authenticated. Note that the system notifies the user after authentication that user actions are being recorded.

To record a trace file:

1. In the admin console, choose Maintenance > Troubleshooting > User Sessions > Session Recording.

[Session Recording](#) shows the session recording page.

2. Enter the username of the user whose session you want to record.
3. Select the Web (DSRecord) check box to record the user's web session and then select the Ignore browser cache check box if you want to ignore cached copies of the problem web site, which the system would not otherwise record as a part of the trace file (optional).
4. Select the Client/Server (for JCP) check box to record Java Communication Protocol client/server application sessions (optional).
5. Click Start Recording. The system signs out the user.
6. Instruct the user to sign in again and browse to the problem web site or connect to the client/server application.
7. Click Stop Recording.
8. Download the trace file(s) from the Current Trace File section:
 - Click the **DSRecord Log** link to download the Web trace file.
 - Click the **JCP or NCP Client-Side Log** link to download the client/server application trace file.
9. E-mail the file(s) to Ivanti Support for review.

Session Recording

Troubleshooting

User Sessions | Monitoring | Tools | System Snapshot | Remote Debugging

Policy Tracing | Simulation | **Session Recording** | Virtual Desktop

Occasionally, users may encounter a Web site that is not displayed properly when viewed through the IVE. Det generated while browsing the problem site. This page allows you to record such a trace for a given user.

After the user has browsed the problem site, you should turn recording off, inspect the trace file, and then send trace will force a signed-in user to sign in again.

Record Trace File	Current Trace File
Status: Not recording User: <input type="text"/> Realm: (Select a realm) ▾ <input checked="" type="checkbox"/> Web (DSRecord) <input checked="" type="checkbox"/> Ignore browser cache <input checked="" type="checkbox"/> Client/Server (for Java clients) <input type="button" value="Start Recording"/>	Trace files: None

Using the Debug Log

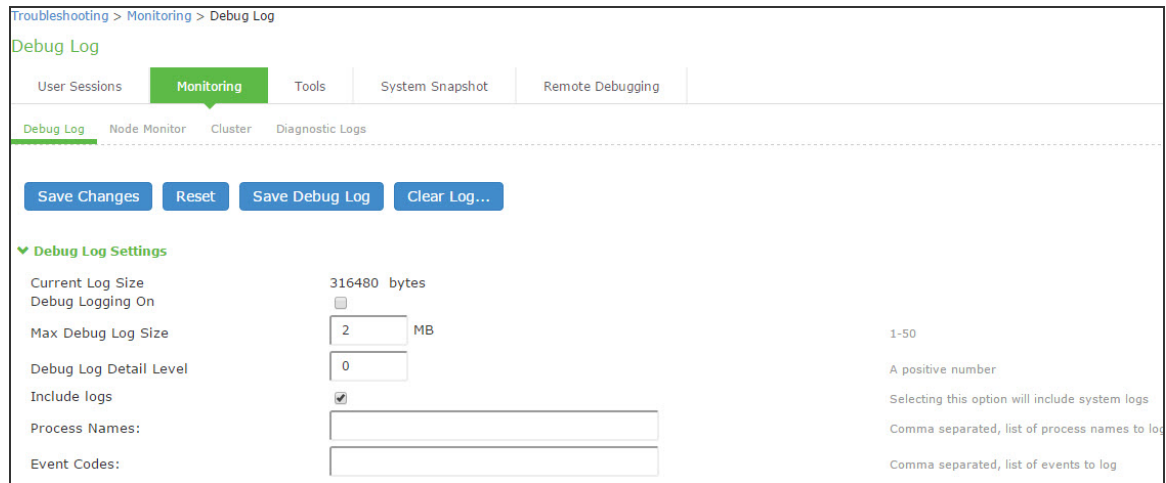
The Global Support Center might direct you to create a debug log to assist them in helping you debug an issue with the system. The debug log is used only by Global Support Center.

To use debug logging:

1. Select **Troubleshooting > Monitoring > Debug Log** to display the configuration page.
[Debug Logging Configuration Page - Ivanti Connect Secure](#) shows the configuration page for Ivanti Connect Secure.
2. Complete the configuration as described in [Table](#).
3. Click **Save Changes**. When you save changes with Debug Logging On selected, the system begins generating debug log entries.
4. Initiate the action you want to debug, such as a user sign in. You can reset the debug log file to restart debug logging if it takes you too long to initiate the action.

5. Click **Save Debug Log** to save the debug log to a file that you can send to Global Support Center. You can clear the log after you have saved it to a file.
6. Unselect Debug Logging On and click **Save Changes** to turn off debug logging.

Debug Logging Configuration Page - Ivanti Connect Secure



The following table lists the Debug Log Configuration Guidelines:

Settings	Guidelines
Current Log Size	Displays the size of the current log file. If it is large, use the controls to save, reset, or clear the log file.
Debug Logging On	Select to turn on debug logging.
Debug Log Size	Specify a maximum debug log file size. The default is 2 MB. The maximum is 250 MB.
Debug Log Detail Level	Specify the debug log detail level. Obtain this from Global Support Center.
Include logs	Select this option to include system logs in the debug log file. Recommended.
Process Names	Specify the process name. Obtain this from Global Support Center.
Event Codes	Specify the event code. Obtain this from Global Support Center.

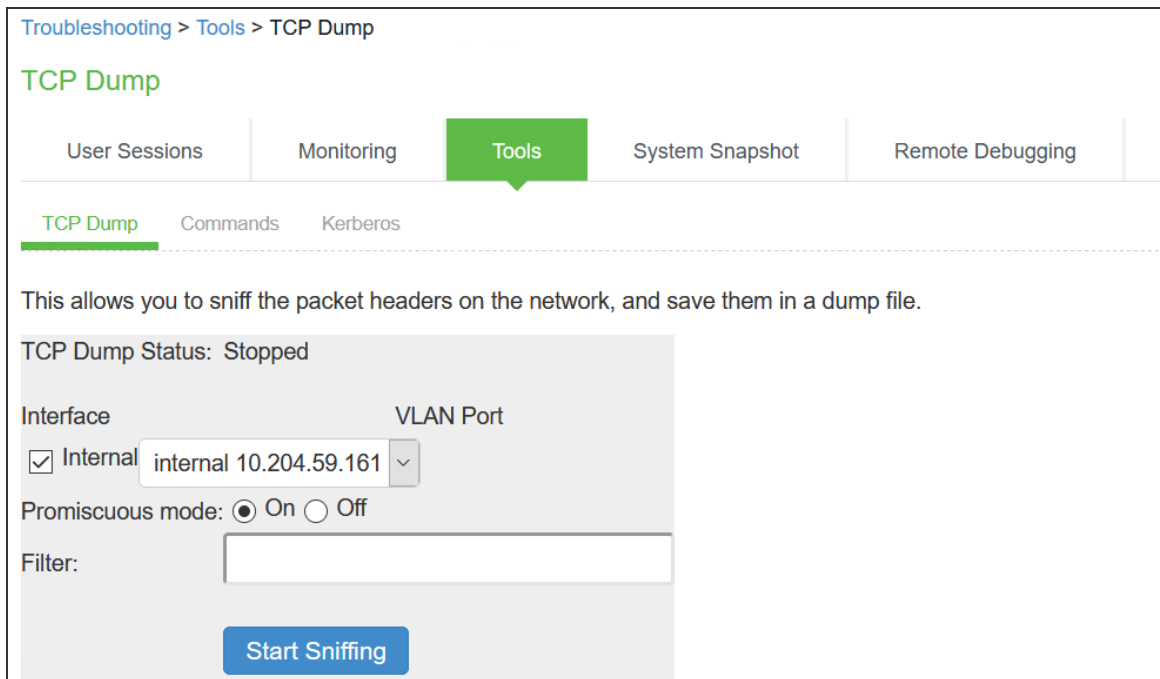
Using the tcpdump Utility

You can run the tcpdump utility from the admin console.

To use tcpdump:

1. Select **Troubleshooting > Tools > TCP Dump** to display the configuration page.
[TCP Dump Configuration Page - Ivanti Connect Secure](#) shows the configuration page for Ivanti Connect Secure.
2. Complete the configuration as described in [Table](#).
3. Click **Start Sniffing** to start the tcpdump process.
4. Initiate the action you want to debug, such as a user sign in.
5. Click **Stop Sniffing** to write the tcpdump output to the screen.
6. Click **Get** to save the output to a file, or click Delete to clear the output.

TCP Dump Configuration Page - Ivanti Connect Secure



The following table lists the Debug Log Configuration Guidelines:

Settings	Guidelines	
TCP Dump Status	Displays whether the utility is stopped or running.	
Interface	Select the ports on which to sniff.	
VLAN Port	Select the VLAN port.	
Promiscuous mode	Select a promiscuous mode option.	
Filter	Specify a filter expression. For information about TCP dump filter expressions, see the UNIX man page.	
	Example	Result
	tcp port 80	Sniffs packets on TCP port 80.
	port 80	Sniffs packets on TCP or UDP port 80.
	ip	Sniffs the IP protocol.
	tcp	Sniffs the TCP protocol.
	dst #.#.#.#	Sniffs the destination IP address specified, where #.#.#.# is a valid IP address.
	src #.#.#.#	Sniffs the source IP address specified, where #.#.#.# is a valid IP address.
	port 80 or port 443	Sniffs on port 80 or port 443.
	src #.#.#.# and dst #.#.#.#	Sniffs the source and destination IP addresses or hosts specified, where each #.#.#.# represents a valid IP address.
tcp port 80 or port 443 and dst #.#.#.# and src #.#.#.#	This example shows how to specify multiple parameters to create a filter that sniffs on TCP port 80, or on TCP or UDP port 443, and on the destination and source ports, where each #.#.#.# represents a valid IP address.	

Using the Samba Diagnostic Log

The Samba diagnostic log utility allows you to view trace and debug the samba troubleshooting messages on the new AD authentication server. When samba diagnostic logging is enabled, the internal logs related to AD authentication server is generated.

Observe the following guidelines:

- Diagnostic logging affects system performance.
- Must be used only when the admin UI error messages, event logs and admin logs are not very useful.
- Enabling/Disabling samba logs will restart certain modules and user logins may fail during the restart.
- The default debug log setting will generate minimal logs. Enabling debug log with event AAA or AAA::samba along with this feature can generate more logs based on the debug log level.
- Enabling samba logs will cause logs to be generated from all configured AD authentication servers. Logs from multiple AD servers are interleaved and can be identified by the header in each line of the logs.

To use samba diagnostic logging:


1. Select **Troubleshooting > Monitoring > Diagnostic Logs** to display the configuration page.

[Samba Diagnostic Logging Configuration Page - Ivanti Connect Secure](#) shows the configuration page.

2. Complete the configuration as described in [Table](#).
3. Click **Save Changes**. When you save changes with Samba Diagnostic Logging On selected, the system begins generating diagnostic log entries.
4. Initiate the action you want to debug, such as a user sign in.
5. Manage the resulting log:
 - Click **Save Log** to save the log files in a zipped format.
 - Click **Clear Log** to remove previous logs and start diagnostic logging with a fresh file.
 - Click **Save And Clear Log** to save the diagnostic log to a file that you can send to Global Support Center. The existing logs in the device will be cleared after saving.

- Unselect **Samba Diagnostic Logging On** and click **Save Changes** to turn off diagnostic logging.

Samba Diagnostic Logging Configuration Page - Ivanti Connect Secure



The following table lists the Samba Debug Log Configuration Guidelines:

Settings	Guidelines
Samba Diagnostic Logging On	Select this option to generate logs related to AD server.
Max Diagnostic Log Size	Specify a maximum log file size between 1 to 100 MB. Default log size is 10 MB.

Using the SNMP Diagnostic Log

This SNMP diagnostic log utility is used for generating debug logs related SNMP daemon. This utility is available only when the **SNMP Queries** and **SNMP Traps** options are enabled in the System > Log/Monitoring > SNMP page.

This option will be automatically turned off after 15 minutes (or) when diagnostic log reaches 10 MB (whichever is earlier).

To use SNMP diagnostic logging:

1. Select **Troubleshooting > Monitoring > Diagnostic Logs** to display the configuration page.
[Samba Diagnostic Logging Configuration Page - Ivanti Connect Secure](#) shows the configuration page.
2. In the SNMP Diagnostics Log Settings section, select the **SNMP Diagnostic Logging On** check box.
3. Click **Save Changes**.

Using the REST Monitor


With the REST Monitoring tool, administrator can enable REST based monitoring of the ICS device. When client makes REST call using HAWK authentication and credentials, ICS sends the information about CPU, memory and load average.

To enable / disable REST monitoring:

1. Select **Troubleshooting > Monitoring > REST Monitor** to display the configuration page.
Figure shows the configuration page for Monitoring ICS over REST.
2. Click **Enable** to activate REST monitoring.
3. Click **New Client** to add client.
4. In the Create Client page displayed, enter a unique **Client Name** to identify the client and applicable Password.
5. Click **Save Changes**.
6. To modify a client name, click the corresponding client name link.
7. To delete a client, select the corresponding check box and click **Delete**.
8. To disable monitoring, click **Disable**.

When REST monitoring is enabled or disabled, the information is logged under Admin logs.

REST Monitoring Configuration Page (Main Menu: **Maintenance**)



Using Network Troubleshooting Commands

You can run common network troubleshooting commands such as arp, ping, ping6, traceroute, traceroute6, NSlookup, and AvgRTTs from the admin console. You can use these connectivity tools to see the network path from the system to a specified server. If a client can ping or traceroute to the access system, and the access system can ping the target server, any remote users should be able to access the server through the access system.

To run network troubleshooting commands:

1. Select **Troubleshooting > Tools > TCP Commands** to display the configuration page.

[Network Troubleshooting Commands Configuration Page - Ivanti Connect Secure](#) shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration as described in [Table](#).
3. Click **OK** to run the command and write the output to the screen.
4. Click **Clear** to clear the output.

[Network Troubleshooting Commands Configuration Page - Ivanti Connect Secure](#)

Troubleshooting > Tools > Commands

Commands

User Sessions | Monitoring | **Tools** | System Snapshot | Remote Debugging

TCP Dump | **Commands** | Kerberos | Licensing Protocol Trace

Command: NSLookup

Query Type: Ping
Ping6
Traceroute
Traceroute6
NSLookup
ARP

Query: Information

DNS Server: Hostname or IP address of specific DNS server to use (optional)

Interface: AvgRTTs
Portprobe

VLAN Port: Internal Port (10.96.3.3)

Internal Port Management Port

OK Clear

The following table lists the Network Troubleshooting Commands Configuration Guidelines:

Settings	Guidelines
Command	<p>Select a network troubleshooting command:</p> <p>Ping/Ping6-Use the ping command to verify that the system can connect to other systems on the network. In the event of a network failure between the local and remote nodes, you do not receive a reply from a pinged device. In that case, contact your LAN administrator for help. The ping command sends packets to a server and returns the server response, typically a set of statistics including the target server's IP address, the time spent sending packets and receiving the response, and other data. You can ping unicast or multicast addresses, and you must include the target server name in the request. Select ping to ping an IPv4 address or hostname. Select ping6 to ping an IPv6 address or hostname.</p> <p>Traceroute/Traceroute6-Use the traceroute command to discover the path that a packet takes from Ivanti Connect Secure to another host. Traceroute sends a packet to a destination server and receives an ICMP TIME_EXCEEDED response from each gateway along its path. The TIME_EXCEEDED responses and other data are recorded and displayed in the output, showing the path of the packet round-trip. Select traceroute to target an IPv4 address or hostname. Select traceroute6 to target an IPv6 address or hostname.</p> <p>NSlookup-Use NSlookup to get detailed information about a name server on the network. You can query on several different types of information, including a server's IP address, alias IP address, start-of-authority record, mail exchange record, user information, well-known services information, and other types of information.</p> <p>ARP-Use the arp command to map IP network addresses to the hardware addresses. The Address Resolution Protocol (ARP) allows you to resolve hardware addresses. To resolve the address of a server in your network, a system sends information about its unique identifier to a server process executed on a server in the intranet. The server process then returns the required address to the client process.</p> <p>AvgRTTs-Use AvgRTTs to display the average round-trip time (RTT) to the localhost.</p> <p>Portprobe-Display the Transmission Control Protocol (TCP) or the User Datagram Protocol (UDP) port status (open or closed).</p>
Target server	Specify the IP address or hostname for the target server.
Interface	Select the interface from which to send the command.

Settings	Guidelines
VLAN Port	Select the VLAN through which the connectivity needs to be checked.
Output	Displays command output.

Troubleshooting TCP and UDP Port Status

Problem	Description: The system makes several connections to back-end servers using various port numbers. If communication between the system and the back-end servers stops, it can be difficult to determine the source of the problem.
Solution	You can use the Portprobe command to display the Transmission Control Protocol (TCP) or the User Datagram Protocol (UDP) port status (open or closed).



Only the system internal ports, management port and internal VLAN ports support the Portprobe command.

A TCP port can be closed under two conditions:

- The system sends a connection request to the back-end server port and the back-end server closes the connection (sends an RST packet).
- The connection request times out because the back-end server is not found, or the back-end server is too busy to respond to the connection request.

If either of these conditions occurs, the system sends a ping command to the back-end server. If the ping command is successful, the back-end server is considered reachable, but the back-end server port is closed. If the ping command fails, the back-end server is considered unreachable.

For UDP ports, the system sends a UDP datagram with a ping to the back-end server port. If the back-end server responds with Internet Control Message Protocol (ICMP) port unreachable or ICMP unreachable, the back-end port is considered unreachable. If the back-end server responds with ICMP host unreachable then the back-end server is considered unreachable.

To troubleshoot the TCP or UDP port:

1. Select **Maintenance > Troubleshooting > Tools > Commands**.
2. Select the Portprobe command.
3. Select either TCP or UDP.

4. Enter the target server and port number. You can enter an IP address, hostname or FQDN for the target server.
5. Enter the probe count. This is the number of times the system attempts to communicate with the back-end server port. The default for TCP is one; the default for UDP is five.
6. Enter the probe timeout. This is the number of seconds the system waits for a response from the back-end server port.
7. Select either the internal port or the management port. If the management port is not configured, it is not displayed.
8. If using an internal port, select the internal VLAN port from the list.
9. Click **OK**.

[Successful TCP Port Probe](#) show an example of a successful and an unsuccessful port probe.

Successful TCP Port Probe

Troubleshooting > Tools > Commands

Commands

User Sessions | Monitoring | **Tools** | System Snapshot | Remote Debugging

TCP Dump | **Commands** | Kerberos | Licensing Protocol Trace

IPv6 is not supported

Command:

Protocol:

Target Server: Target port[1-65535]:

Probe Count: (default: tcp = 1, udp = 5, max = 100)

Probe Timeout: (default: 1 secs, max = 180 secs)

Interface: Internal Port Management Port

VLAN Port:

Output:

```
Resolving IP address for www.google.com
Resolved IP address: 172.217.18.68
Starting port probing

Tcp probe : 172.217.18.68:80 Open
Tcp probe : 172.217.18.68:80 Open
Tcp probe : 172.217.18.68:80 Open
Tcp probe : 172.217.18.68:80 Open
Tcp probe : 172.217.18.68:80 Open
Tcp probe : 172.217.18.68:80 Open

Operation complete
```

Unsuccessful UDP Port Probe

Troubleshooting > Tools > Commands

Commands

User Sessions | Monitoring | **Tools** | System Snapshot | Remote Debugging

TCP Dump | **Commands** | Kerberos | Licensing Protocol Trace

IPv6 is not supported

Command:

Protocol:

Target Server: Target port[1-65535]:

Probe Count: (default: tcp = 1, udp = 5, max = 100)

Probe Timeout: (default: 1 secs, max = 180 secs)

Interface: Internal Port Management Port

VLAN Port:

Output:

```

Resolving IP address for 10.209.118.10
Resolved IP address: 10.209.118.10
Starting port probing

UDP probe: 10.209.118.10:8888 Close ( Destination host is unreachable )
UDP probe: 10.209.118.10:8888 Close ( Destination host is unreachable )
UDP probe: 10.209.118.10:8888 Close ( Destination host is unreachable )
UDP probe: 10.209.118.10:8888 Close ( Destination host is unreachable )
UDP probe: 10.209.118.10:8888 Close ( Destination host is unreachable )

Operation complete
    
```

Running NSLookup to Test Name Server Connectivity

To run NSLookup to test name server connectivity:

1. In the admin console, choose **Maintenance > Troubleshooting > Tools > Commands**.

[Network Troubleshooting Commands Configuration Page - Ivanti Connect Secure](#) shows the configuration page for Ivanti Connect Secure.

2. From the Command list, select NSLookup.

3. Select the type of query to use from the Query Type drop down menu.
4. Enter the query, which is a hostname, an IP address, or other information, depending on your selection of query type.
5. Enter the DNS server name or IP address.
6. Enter other options.
7. Click **OK** to run the command.

Network Troubleshooting Commands Configuration Page - Ivanti Connect Secure

The screenshot displays the 'Commands' configuration page. The breadcrumb trail is 'Troubleshooting > Tools > Commands'. The 'Tools' tab is active, and the 'Commands' sub-tab is selected. The form contains the following fields:

- Command:** A dropdown menu with 'NSLookup' selected.
- Query Type:** A dropdown menu with options: Ping, Ping6, Traceroute, Traceroute6, and NSLookup (highlighted).
- Query:** A text input field with a placeholder: 'Hostname, or IP address, or other information based on Query Type'.
- DNS Server:** A text input field with a placeholder: 'Hostname or IP address of specific DNS server to use (optional)'.
- Interface:** Radio buttons for 'Internal Port' and 'Management Port'.
- VLAN Port:** A dropdown menu with 'Internal Port (10.96.3.3)' selected.

At the bottom of the form are 'OK' and 'Clear' buttons.

Using the Kerberos Debugging Utility

You can run the Kerberos debugging utility from the admin console. The utility checks the DNS infrastructure for validity of the Kerberos realms and defined credentials.

To use the Kerberos debugging utility:

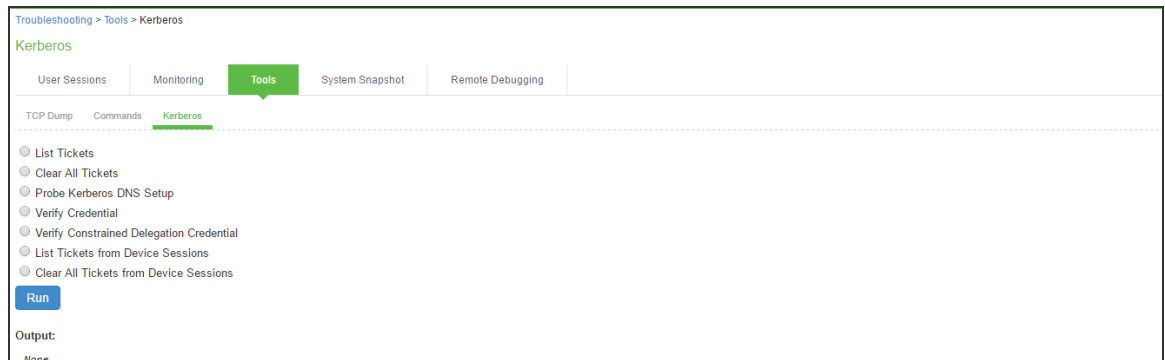
1. Select **Maintenance > Troubleshooting > Tools > Kerberos** to display the configuration page.

[Kerberos Debugging Utility Configuration Page](#) shows the configuration page.

2. Complete the configuration as described in [Table](#).
3. Click **Run** to start the debugging process.

4. Click **Get to save the output to a file**, or click **Delete** to clear the output.

Kerberos Debugging Utility Configuration Page



The following table lists the Kerberos Debugging Utility Configuration Guidelines:

Settings	Guidelines
List Tickets	Select this option to list all tickets. Specify the username and the realm name.
Clear All Tickets	Select this option to remove all tickets. Specify the username and realm name.
Probe Kerberos DNS Setup	Select this option to display the configuration elements for the Kerberos DNS test. Specify the realm name and the fully qualified domain name.
Verify Credential	Select this option to verify the Kerberos ticket is valid. Specify the following: Kerberos Client Server Client Realm Server Realm (Optional) Client KDC Server KDC (Optional) Password For example, if you use Kerberos to verify the username and password provided by the user, this option verifies the credentials it obtains to make sure they belong to a trusted KDB site.

Settings	Guidelines
Verify Constrained Delegation Credential	Select this option to verify the Constrained Delegation ticket is valid. Specify the following: Kerberos Client Delegation Account Server Client Realm Server Realm (Optional) Client KDC Server KDC (Optional) Password
List Tickets from Device Sessions	Select this option to list all tickets from device sessions. Specify the following: Username
Clear All Tickets from Device Sessions	Select this option to clear all tickets from device sessions. Specify the following: Username
Output	Displays results of the probe, for example: KDCs for realm matrix.net: top.matrix.net, top.matrix.net Operation complete

Using System Snapshots

A snapshot of the system state captures details that can help Global Support Center diagnose system performance problems. The system stores up to ten snapshots, which are packaged into an encrypted "dump" file that you can download and then e-mail to Global Support Center.

To create and manage system snapshots:

1. Select **Maintenance > Troubleshooting > System Snapshot** to display the configuration page.

[System Snapshot Configuration Page - Ivanti Connect Secure](#) shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration and actions as described in [Table](#).

[System Snapshot Configuration Page - Ivanti Connect Secure](#)

troubleshooting > System Snapshot

System Snapshot

User Sessions | Monitoring | Tools | **System Snapshot** | Remote Debugging

A snapshot of the system state captures details that can help Pulse Secure Support diagnose system performance problems. The system stores up to ten snapshots, which are packaged into an encrypted "dump" file that you can download to a network machine and then email to Pulse Secure Support.

Take Snapshot | Take Snapshot on All Nodes | Delete

10 records per page

Snapshot	Size	Date
Periodic snapshot(with debuglog, config)	1048275 bytes	2016-03-24 10:13:18
Periodic snapshot(with debuglog, config)	1039064 bytes	2016-03-24 06:13:18
Periodic snapshot(with debuglog, config)	1026893 bytes	2016-03-24 02:13:18
Periodic snapshot(with debuglog, config)	1028702 bytes	2016-03-23 22:13:17
Periodic snapshot(with debuglog, config)	1015052 bytes	2016-03-23 18:13:17
Periodic snapshot(with debuglog, config)	987940 bytes	2016-03-23 14:13:17
Periodic snapshot(with debuglog, config)	974102 bytes	2016-03-23 10:13:17
Periodic snapshot(with debuglog, config)	952022 bytes	2016-03-23 06:13:18
Periodic snapshot(with debuglog, config)	947571 bytes	2016-03-23 02:13:17
Periodic snapshot(with debuglog, config)	957126 bytes	2016-03-22 22:13:17

System snapshot options

- Include system config
- Include debug log
- Schedule automatic snapshots

You should enable automatic scheduled snapshots only when asked to do so by Pulse Secure Support as part of a troubleshooting operation. Enabling this feature can affect system performance. In most situations, a four-hour snapshot schedule captures the needed data without impacting system performance. Do not set a schedule interval of less than 30 minutes as this can affect system performance.

Take a snapshot every: 12 hours (0 - 336 hours)

0 minutes (0 - 59 minutes)

Max size allocated for periodic snapshots: 60 MB (60 - 500 MB)

Stop taking snapshots after: Date: (mm/dd/yyyy)

Time: (hh:mm) AM

Disable debug logs at stop time:

Save Changes

The following table lists the System Snapshot Configuration Guidelines:

Settings	Guidelines
Include system config	Include the system configuration file in the snapshot.
Include debug log	Include debug logs (if any).

Settings	Guidelines
Schedule Automatic Snapshots	Enable automatic scheduled snapshots only when asked to do so by Ivanti Support as part of a troubleshooting operation. Enabling this feature can affect system performance. In most situations, a four-hour snapshot schedule captures the needed data without impacting system performance. Do not set a schedule interval of less than 30 minutes as this can affect system performance
	Frequency-Specify a frequency in hours and minutes.
	Maximum size-Specify a maximum file size.
	Stop taking snapshots-Specify a date and time to stop the automatic snapshot job.
	Disable debug logs at stop time-Specify that you also want to turn off debug logging when you stop the automatic snapshot job.
Save	Save the configuration.
Take Snapshot	Generate a snapshot now.
Delete	Delete a snapshot file.

Using Remote Debugging

Remote debugging allows Global Support Center to directly access this system over a secure connection. You should enable this feature only if you have been requested to do so by Global Support Center in response to an issue that you have reported.

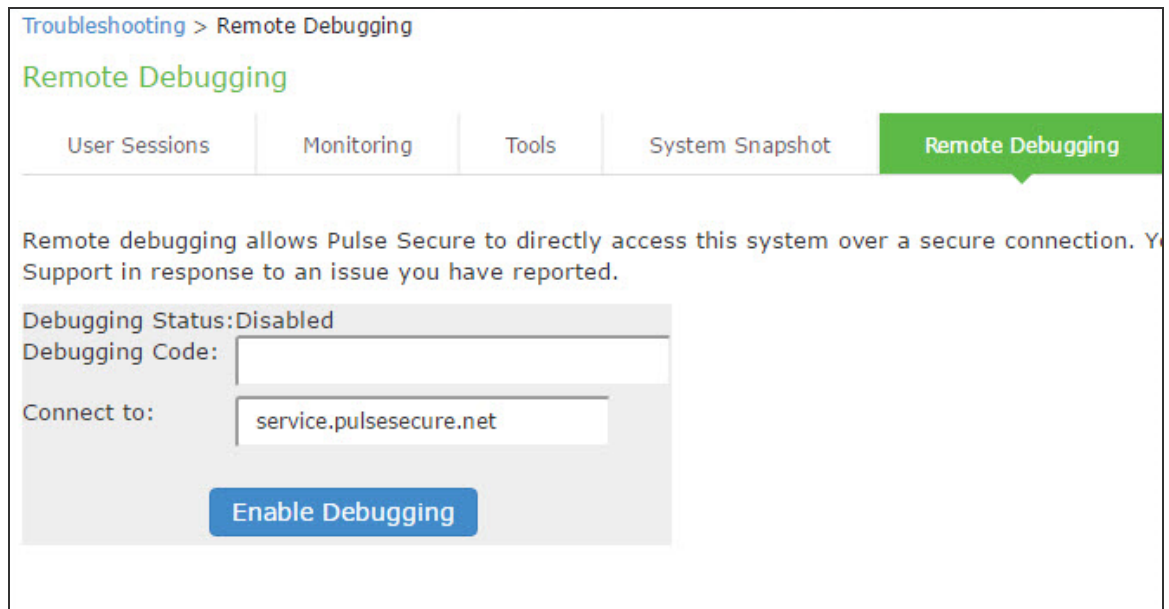
To enable remote debugging:

1. Select **Maintenance > Troubleshooting > Remote Debugging** to display the configuration page.

[Remote Debugging Configuration Page - Ivanti Connect Secure](#) shows the configuration page for Ivanti Connect Secure.

2. Complete the configuration and actions as described in [Table](#).

[Remote Debugging Configuration Page - Ivanti Connect Secure](#)



The following table lists the Remote Debugging Configuration and Action Guidelines:

Settings	Guidelines
Debugging Status	Displays whether remote debugging is enabled or disabled.
Debugging Code	Specify a code as instructed by Global Support Center.
Connect to	Specify the fully qualified domain name as instructed by Global Support Center.
Enable Debugging	Click this option to allow remote debugging.

Using Log Selection

The various system logs and troubleshooting logs that help in investigating user access issues and system issues can be configured and accessed from the Log Selection page.

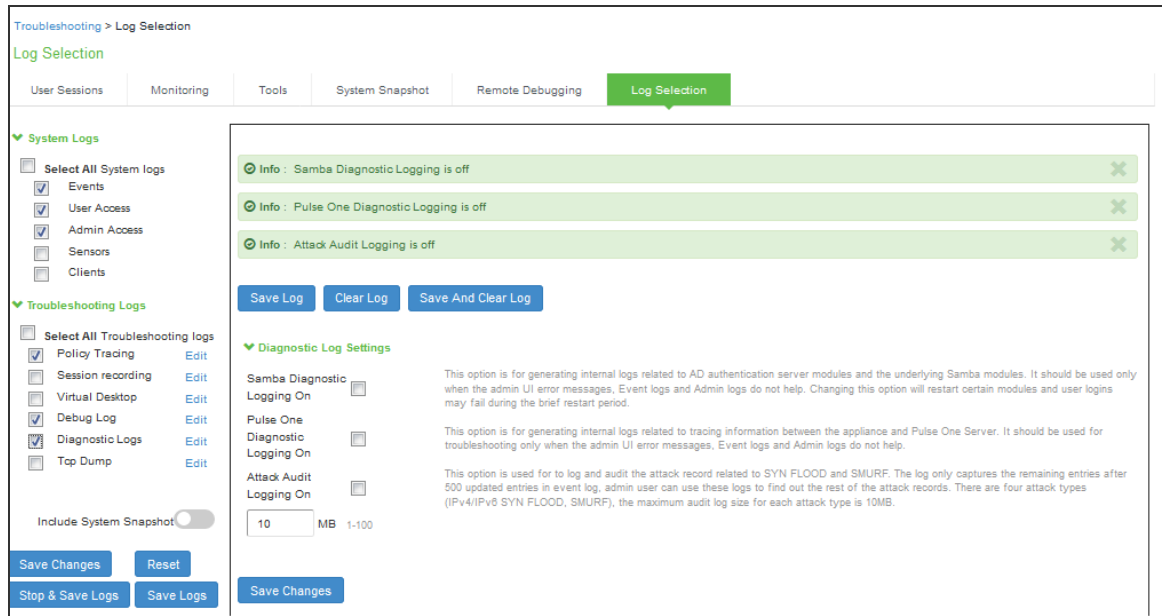
To configure system logs and troubleshooting logs:

1. Select **Maintenance > Troubleshooting > Log Selection** to display the Log Selection page.

[Log Selection Page](#) shows the Log Selection page.

2. Complete the configuration and actions as described in [Table](#).

Log Selection Page




The following table lists the Log Selection Configuration Guidelines:

Settings	Guidelines
System Logs	
Select All System Logs	Select this check box to capture all system logs. To choose specific log, select individual system log from the list.
Troubleshooting Logs	
Select All Troubleshooting Logs	Select this check box to capture all troubleshooting logs. To choose specific log, select individual troubleshooting log from the list.
Edit log settings	To configure the settings of individual logs, click the corresponding Edit link. Complete the configuration and click Save Changes.
Stop and Save Logs	Stops the services used for the log collection and archives all the selected logs and then prompts to download the archive file.
Save Logs	Archives all the selected logs and prompts to download it as a bundle.

Clustering

Clusters define a collection of servers that operate as if they were a single machine. A cluster pair is used to refer to a cluster of two units and a multiunit cluster refers to a cluster of more than two units. Once two or more units are joined in a cluster, they act as one unit.

The following figure shows two PSA series devices deployed as a cluster pair:



Ivanti access management framework supports two types of clusters:

- Load balancing clusters or active/active clusters
- Failover clusters or active/passive clusters

For details about the configuration, various deployment scenarios, reports, etc. refer to the *ICS Clustering Configuration Guide* available at <https://www.ivanti.com/support/product-documentation>.

Delegating Administrator Roles

About Delegating Administrator Roles

The access management system enables you to delegate various management tasks to different administrators through system administrator roles and security administrator roles. System and security administrator roles are defined entities that specify management functions and session properties for administrators who are mapped to those roles. You can customize an administrator role by selecting the feature sets, user roles, authentication realms, resource policies, and resource profiles that members of the administrator role are allowed to view and manage. Note that system administrators may only manage user roles, realms, and resource policies; only security administrators can manage administrator components.

For example, you can create a system administrator role called "Help Desk Administrators" and assign users to this role who are responsible for fielding tier 1 support calls, such as helping users understand why they cannot access a Web application or system page. In order to help with troubleshooting, you may configure settings for the "Help Desk Administrators" role as follows:

- Allow the help desk administrators Write access to the System > Log/Monitoring page so they can view and filter the system logs, tracking down critical events in individual users' session histories, as well as the Maintenance > Troubleshooting page so they can trace problems on individual users' systems.
- Allow the help desk administrators Read access to the Users > User Roles pages so they can understand which bookmarks, shares, and applications are available to individual users' roles, as well as the Resource Policy or Resource Profile pages so they can view the policies that may be denying individual users access to their bookmarks, shares, and applications.
- Deny the help desk administrators any access to the remaining System pages and Maintenance pages, which are primarily used for configuring system-wide settings-such as installing licenses and service packages-not for troubleshooting individual users' problems.



In addition to any delegated administrator roles that you may create, the system also includes two basic types of administrators: super administrators (.Administrators role), who can perform any administration task through the admin console and read-only administrators (.Read-only Administrators role), who can view-but not change-the entire system configuration through the admin console.

You can also create a security administrator role called "Help Desk Manager" and assign users to this role who are responsible for managing the Help Desk Administrators. You might configure settings for the "Help Desk Manager" role to allow the Help Desk Manager to create and delete administrator roles on his own. The Help Desk Manager might create administrator roles that segment responsibilities by functional areas of the system. For example, one administrator role might be responsible for all log monitoring issues. Another might be responsible for all Network Connect problems.

All devices allow members of the .Administrators role to configure general role settings, access management options, and session options for the .Administrators and .Read-Only Administrators roles.



On certain pages, such as the role mapping page, the delegated administrator can view the role names even though the administrator does not have read/write access. However, the delegated administrator cannot view the details of that role.

Creating and Configuring Administrator Roles

You can use the Administrators > Admin Roles pages to set default session and user interface options for delegated administrator roles.

To create individual administrator accounts, you must add the users through the appropriate authentication server (not the role). For example, to create an individual administrator account, you may use settings in the Authentication > Auth. Servers > Administrators > Users page of the admin console. For detailed instructions on how to create users on the Administrators server and other local authentication servers. For instructions on how to create users on third-party servers, see the documentation that comes with that product.

To create an administrator role:

1. In the admin console, choose **Administrators > Admin Roles**.
2. Do one of the following:
 - Click **New Role** to create a new administrator role with the default settings.
 - Select the check box next to an existing administrator role and click **Duplicate** to copy the role and its custom permissions. Note that you cannot duplicate the system default roles (.Administrators and .Read-Only Administrators).
3. Enter a name (required) and description (optional) for the new role and click **Save Changes**.
4. Modify restrictions, session options, and UI options according to your requirements.



If you select one of the system's default administrator roles (.Administrators or .Read-Only Administrators), you can only modify settings in the General tab (since the default system administrators roles always have access to the functions defined through the System, Users, Administrators, and Resource Policies tabs).

You cannot delete the Administrators and Read Only Administrators roles since they are default roles.

Specifying Management Tasks to Delegate

This topic contains information about delegating management tasks to various delegated administrator roles.

Delegating System Management Tasks

Use the **Administrators > Admin Roles > Select Role > System** tab to delegate various system management tasks to different administrator roles. When delegating privileges, note that:

- The system allows all administrators read-access (at minimum) to the admin console home page (System > Status > Overview), regardless of the privilege level you choose.
- The system does not allow delegated administrators write-access to pages where they can change their own privileges. Only those administrator roles that come with the system (.Administrators and .Read-Only Administrators) may access these pages:
 - **Maintenance > Import/Export (Within this page, .Read-Only Administrators can export settings, but cannot import them.)**
 - **Maintenance > Push Config**
 - **Maintenance > Archiving > Local Backups**

Delegating User and Role Management

Use the **Administrators > Admin Roles > Select Role > Users > Roles** sub-tab to specify which user roles the administrator role can manage. When delegating role management privileges, note that:

- Delegated administrators can only manage user roles.
- Delegated administrators cannot create new user roles, copy existing roles, or delete existing roles.

- If you allow the delegated administrator to read or write to any feature within a user role, the system also grants the delegated administrator read access to the Users > User Roles > Select Role > General > Overview page for that role.
- If you grant a delegated administrator write access to a resource policy through the Administrators > Admin Roles > Select Administrator Role > Resource Policies page, he may create a resource policy that applies to any user role, even if you do not grant him read access to the role.

Delegating User Realm Management

Use the Administrators > Admin Roles > Select Role > Users > Authentication Realms tab to specify which user authentication realms the administrator role can manage. When delegating realm management privileges, note that:

- System administrators can only manage user realms.
- System administrators cannot create new user realms, copy existing realms, or delete existing realms.
- If you allow the system administrator to read or write to any user realm page, the system also grants the system administrator read-access to the Users > User Realms > Select Realm > General page for that role.

Delegating Administrative Management

Use the Administrators > Admin Roles > Select Roles > Administrators tab to specify which system administrator roles and realms the security administrator role can manage. When delegating security administrative privileges, note that:

- The security administrator role provides control over all administrative roles and realms.
- You can give a security administrator control exclusively over administrator roles, over administrator realms, or over both.
- You can restrict or grant the security administrator the permission to add and delete administrator roles and administrator realms.

Delegating Resource Policy Management

Use the Administrators > Admin Roles > Resource Policies tab to specify which user resource policies the administrator role can manage. When delegating resource policy management privileges, note that delegated system administrators cannot modify the following characteristics of resource policies:

- The resource itself (that is, the IP address or hostname).
- The order to evaluate the resource policies.

Delegating Resource Profile Management

Use the Administrators > Admin Roles > Resource Profiles tab to specify which user resource profiles the administrator role can manage. When delegating resource profile management privileges, note that delegated system administrators cannot modify the following characteristics of resource profiles:

- The resource itself (that is, the IP address or hostname)
- The order to evaluate the resource policies.

Deployments with IDP

About IDP

Securing intranet work application and resource traffic is vital to protecting your network from hostile outside intrusion. You can add levels of application security to your remote access network by integrating a Ivanti Connect Secure system with a Juniper Networks Intrusion Detection. The IDP device may provide the following types of protection in this solution (some forms of protection depend upon the specific configuration):

The IDP device provides the following types of protection (some forms of protection depend upon the specific configuration):

- Protects against attacks from user to application and from application to user (from a server-side endpoint)
- Detects and blocks most network worms based on software vulnerabilities
- Detects and blocks non-file-based Trojan Horses
- Detects and blocks effects of spyware, adware, and key loggers
- Detects and blocks many types of malware
- Detects and blocks zero-day attacks through the use of anomaly detection



From 9.1R15 onwards, support for Sensors is deprecated. Ensure you remove all related configurations before upgrading to 9.1R15. **Upgrade may fail if all configurations are not removed.** For more information refer [KB45044](#).

You do not need a special license from Ivanti to enable interaction between Ivanti Connect Secure and the IDP.

Using the Ivanti Connect Secure admin console, you can configure and manage interaction attributes between it and an IDP, including the following:

- Global configuration parameters such as the IDP hostname or IP address, the TCP port over which the sensor communicates with Ivanti Connect Secure, and the one-time password Ivanti Connect Secure and IDP use to authenticate with one another.
- Dynamically changing the IDP configuration from Ivanti Connect Secure and alerting the IDP of changes in the IP address pool available to remote users.
- Various levels of attack severity warnings.


The IDP sits behind Ivanti Connect Secure on your internal network and monitors traffic flowing from Ivanti Connect Secure into the LAN. Any abnormal events detected by the IDP Sensor are reported to Ivanti Connect Secure, which you configure to take appropriate action based on the severity level of the reported events.

IDP Deployment Scenarios

The two most likely deployment scenarios are as follows:


- Customer use of Ivanti Connect Secure for extended enterprise access and IDP for security of all perimeter traffic including but not limited to traffic from Ivanti Connect Secure. [Ivanti Connect Secure and IDP Topology Scenario 1](#) illustrates this scenario, in which Ivanti Connect Secure is deployed in the DMZ or on the LAN and the IDP is deployed in-line behind the firewall and in front of the LAN.

Ivanti Connect Secure and IDP Topology Scenario 1



- In the second deployment scenario, IDP is only used to protect traffic that comes through Ivanti Connect Secure but not in-line with other perimeter traffic. Figure 315 illustrates this deployment scenario.

Ivanti Connect Secure and IDP Topology Scenario 2



Configuring Ivanti Connect Secure to Interoperate with IDP

The IDP Sensor is a powerful tool to counter users who initiate attacks. Integration with Ivanti Connect Secure allows you to configure automatic responses as well as manually monitor and manage users.

To configure the system to interoperate with an associated standalone IDP Sensor, you must first ensure the IDP has been configured according to the instructions described in the Signaling Setup appendix of *IDP Series Concepts and Examples Guide, Version 5.1rX*.

On Ivanti Connect Secure you can specify actions to be taken in response to users that perform attacks:

- **Users page**-Manually identify and quarantine or disable users on the System > Status > Active Users page, which lists users who have performed attacks.

Interaction Between the IC Series and IDP

Ivanti Connect Secure receives the source and destination IP addresses and port numbers of the attacking host and the resource against which the attack was launched, along with the attack identifier, severity of the attack, and the time at which the attack was launched.

Ivanti Connect Secure incorporates and displays the attack information received from the IDP sensor on the **System > Status > Active Users** page. Based on the attackers IP address and port number, the system can uniquely identify the user's session.

For manual action, you look up the information available on the Active Users page and decide on an action.

Identifying and Managing Quarantined Users Manually

When the system quarantines a user based on an attack, you can display and manage the states by locating the user link in the **System > Status > Active Users** page.

- A small warning icon displayed in front of the username.
- The hyperlinked username.
- An enabled Quarantined option button on the specific user's page. If the user is not quarantined, the option button is disabled.

You can manage quarantined users from either the admin GUI or by logging in as a user with administrative rights on the local authentication server.

To manage quarantined users:

1. Identify quarantined users at **System > Status > Active Users**.

2. Locate the quarantined user from the **Authentication > Auth. Servers > System Local** on the admin GUI or from the Admin Users window on the local authentication server. You must be logged in to the local authentication server as an administrator user in order to see the Admin User option.
3. Click the username link. The user page opens, showing a number of options. See [Managing Quarantined Users](#)

Managing Quarantined Users

Auth Servers > System Local > Users > Update Local User shan

Update Local User shan

Full Name:

Authenticate using: System Local

Password:

Confirm Password:

One-time use (disable account after the next successful sign-in)

Enabled

Disabled

Quarantined

Require user to change password at next sign in

Note: You must also configure password management on the [Authentication server Settings](#) with 'Allow users to change their passwords' option enabled. Use options on the Administrators/Users > Authentication > [Realm] > Authentication Policy > Password page to specify which realms should inherit the server's password management capabilities.

4. Click **Disabled** to disallow a user from authenticating.
5. Click **Quarantined** to leave a user in a quarantined state. The Quarantined option is only enabled if the user is already quarantined.

i The system assigns quarantined users to the quarantined role, regardless of their login realm.

6. Click **Save Changes**.
7. To re-enable previously quarantined or disabled users, select **Authentication > Auth. Servers > Select Server > Users** and click the link for the given user.

i You can also disable users from this location.

8. Click **Enabled** to release the user from quarantine.
9. Click **Save Changes**.

Dashboard and Reports

A dashboard is an interface used to manage the Ivanti access management framework. It provides an integrated view of all devices and users accessing the network, their device profile information, authentication methods used to gain access, device posture compliance and so on.

A report is an element of a dashboard used to convey complex data in simplified formats. Ivanti access management framework collects log and configuration data from across your network, and it then aggregates the data into reports for you to view and analyze. It provides a standard set of predefined reports that you can use and customize to fit your needs. The reports are grouped into logical categories for information related to authentication, session traffic, device administration, configuration and administration, and troubleshooting.

For details about the configuration, refer to the *ICS Dashboard and Reports Configuration Guide* available at <https://www.ivanti.com/support/product-documentation>.

Pulse One Integration

Overview

ICS appliance can be integrated with the Pulse Workspace console server to auto-provision workspace based on user's group membership and to enable seamless active sync email access for mobile clients. Once this integration is in place, the mobile devices that are managed by Pulse Workspace will get seamless mail access from Enterprise mail server without requiring the users to configure their mail clients.

To configure Pulse Workspace command handlers to auto-provisioning workspace or to enable seamless active sync email access for mobile clients, do the following:

1. Register ICS with Pulse Workspace
2. Maintain Notification Channel
3. Renew Credentials
4. Configure User Role (For seamless Active Sync support)
5. Configure LDAP Authentication Servers to use for Group Lookup (For User's group membership-based auto-provisioning)

Register ICS with Pulse Workspace

ICS has to be registered with Pulse Workspace before it can be used for seamless mail access for Pulse Workspace configured mobile devices. On successful registration, Pulse Workspace sends ICS the following information:

The following table lists the Registration Information:

Registration Information	Description
Hawk Credentials	<p>All communication from ICS to Pulse Workspace are authenticated using the HAWK. Pulse Workspace sends this information in the registration response. The response consists of:</p> <ul style="list-style-type: none"> Key Key Identifier Message Authentication Code Generation Algorithm

Registration Information	Description
Device Identification Information	Each ICS device is uniquely identified in Pulse Workspace. This identification information is sent to ICS in the registration response to be used in all communications.
Notifications Channel URL	To receive any unsolicited notification from Pulse Workspace, ICS creates and maintains a websocket channel with Pulse Workspace. The endpoint URL on the Pulse Workspace for this channel is sent as part of the registration response.
Base API URL	On receiving any unsolicited notification on the websocket, ICS sends a REST request to Pulse Workspace to fetch additional information. The base URL for these REST APIs is sent by Pulse Workspace in the registration response.

Maintain Notification Channel

ICS creates a websocket channel with the Pulse Workspace server. Pulse Workspace sends notification to ICS over this channel. This channel is teared down by the Pulse Workspace once in 24 hours and ICS needs to reconnect to Pulse Workspace on this event. Also, when the HAWK credentials become invalid, the websocket channel is teared down.

ICS keeps the websocket channel up all the time and also takes corrective measures whenever there is a disruption on this channel.

Renew Credentials

HAWK credentials sent by Pulse Workspace are valid for 7 days. After this time, the credentials need to be renewed. When the credentials are in renew state, the notification channel will fail and any communication from ICS to Pulse Workspace cannot be authenticated. The existing credentials can only be used to request the new credentials.

HAWK credentials expire after 30 days. Once the credentials expire, ICS needs to be reconfigured and reregistered using a new registration code. This results into new device identification information and new HAWK credentials.

Configure User Role (For seamless Active Sync support)

Configure the User role that will be used for creating the device records on ICS for Pulse Workspace devices. On creation of a workspace, Pulse Workspace requests ICS to create a device record so that the mobile device which maps to that workspace can access email using ICS as activesync proxy. This requires ICS to know which role should be used for creating the device records. ICS administrator needs to configure this information using the admin UI.

Configure LDAP Authentication Servers to use for Group Lookup (For User's group membership-based auto-provisioning)

Configure the LDAP Authentication server that will be used for handling group validation and user's group membership related requests on ICS for Pulse Workspace Server. ICS administrator needs to configure this information using the admin UI.

Pulse One Configuration

This section covers the configuration required on ICS to enable it to register with the Pulse Workspace console server.

Pulse One Settings (**System > Configuration > Pulse One > Settings**)

Configuration > Pulse One > Settings

Settings

Licensing **Pulse One** Security Certificates DMI Agent NCP Client Types Virtual Desktops User Record Synchronization IKEv2

SAML Mobile VPN Tunneling PSAM Telemetry Advanced Client Configuration Advanced Networking

Settings Command Handlers

*Registration Host: The Host to which the appliance connects to for starting registration flow

*Registration Code: The registration code provided by Pulse One

*Credential Renegotiation Interval: days 1 - 7 days. The time after which credentials are renegotiated

Credentials Exchange time: Fri 2016-03-18 13:16:18 IST The last successful credential exchange time.

▼ Registration Result Details

On successful registration the following information is received from Pulse One

Hashing Algorithm: hs256 Hashing algorithm used for HAWK authentication.

Client Device Id: ec879fe8-418e-4fbb-a636-64580e8b4fd4 Unique id of the appliance on Pulse One

Notification URL: wss://api-jstg.pulseworkspace.net/api/v1/notifications The URL for establishing notification channel

▼ Status Information

Registration Status:

Notification Channel Status:

▼ Actions

The following table lists the Pulse One Configuration Details:

Field	Description
Registration URL	This is the URL to which ICS sends the registration request. The format of the URL is https://<PWS API Host Name>/api/v1/register. The Pulse Workspace API Host name is displayed to the administrator when he/she creates an entry for this appliance on the Pulse Workspace console server.
Registration Code	This is the code that ICS sends to Pulse Workspace in the registration request. This code is generated and displayed to the administrator when he/she creates an entry for this appliance on the Pulse Workspace console server.
Credential Renegotiation Interval	This is the time in days after which ICS automatically does renegotiation of HAWK credentials with Pulse Workspace.
Credentials Exchange time	This is the time at which the last successful credential exchange took place.
Hashing Algorithm	This is the algorithm used for generating the MAC for HAWK authentication. Currently the only supported value is HS256 which is HMAC using SHA-256.
Client Device ID	This is the unique identification information of the ICS device on the Pulse Workspace server. This information is received in the registration response.
Notification URL	This is the URL at which the websocket endpoint is present at the Pulse Workspace server. This information is received in the registration response.
Registration Status	Reports current status of registration. Gray - not yet registered Yellow - registration in progress Green - registered successfully RED - registration failed/renew credentials/credentials expired
Notification Channel Status	Reports current status of notification channel. Gray - not yet connected/connection not required Yellow - connection in progress Green - connected RED - connection failed
Save Changes	Saves the configuration and triggers registration, if required.
Clear configuration	Clears all the configuration and disconnects the notification channel.
Renegotiate credentials	Triggers renegotiation of credentials.



- Hawk is an HTTP authentication scheme providing a method for making authenticated HTTP requests with partial cryptographic verification of the request, covering the HTTP method, request URI, and host.
- To back up and restore Pulse One configuration, administrator should use the binary export/import of system configuration.

Pulse Workspace Handlers

This section covers the configuration of the command handlers that handle the messages received on the notification channel.

Pulse Workspace Handlers (**Administrators > Pulse One**)

The screenshot shows the Pulse One configuration interface. The top navigation bar includes tabs for Licensing, Pulse One (selected), Security, Certificates, DMI Agent, NCP, Client Types, Virtual Desktops, User Record Synchronization, and IKEv2. Below this, there are sub-tabs for SAML, Mobile, and VPN Tunneling. The main content area is titled 'Settings > Command Handlers > Pulse Workspace Handler'. It features a section for 'Active Sync Handler Configuration' with a 'Device Role' dropdown set to '(none)' and a 'Clear Active Sync Device Records' button. To the right of this section is a note: 'Delete all the device records pushed from Pulse Workspace Console Server'. Below this is the 'Group Lookup Handler Configuration' section, which includes a 'Select Authentication Servers to use for group lookup' area. This area has an 'Available Auth Servers' list containing 'LDAP_Server(dc=bluefish,dc=com)', an 'Add ->' button, and a 'Remove' button. The 'Selected Auth Servers' list is currently empty, showing '(none)'. A note below this list states 'Select only one authentication server per domain'. At the bottom left of the configuration area is a 'Save Changes' button.

Active Sync Handler Configuration

This section covers the configuration of the activesync command handlers that create/delete the device records in ICS when Pulse Workspace sends a notification.

The following table describes the Active Sync Handler Configuration:

Field	Description
Device Role	This is the role assigned to the device records created by ICS for the Pulse Workspace registered devices.
Clear Active sync Device Records	This option would delete all the device records pushed from Pulse Workspace Console Server.

- Administrator should ensure that secure email feature is enabled for this user role.
- Use "Clear Active sync Device Records" option only if:
 - This ICS is no longer the active sync provider for Pulse Workspace Server.



- To troubleshoot Device Record sync-up related issues, clear all Pulse Workspace Onboarded Device Records and recreate only the valid Device Records during next active sync Device Record sync-up. Device Record sync-up can happen if there is any new workspace created or existing workspace state is modified or due to periodic sync up initiated by the Pulse Workspace server for every one hour.

Group Lookup Handler Configuration

This section covers the configuration of group lookup command handlers that validate the group existence and also fetches the user's group membership from the configured backend LDAP server when Pulse Workspace sends a notification.

The following table lists the Group Lookup Handler Configuration:

Field	Description
Available Auth Servers	All the configured LDAP Server will be listed under this.
Selected Auth Servers	Select the LDAP authentication server to handle the Group lookup requests.

- Only one authentication server per domain should be selected.
- This functionality is supported only with 'Active Directory' type LDAP server.

- To back up and restore Pulse One command handler configuration, administrator should use the binary export/import of user configuration

Customizable Admin and End-User UIs

Customizable Admin and End-User UIs

The ICS enables you to customize a variety of elements in both the admin console and the end-user interface. This section contains information about which elements you can customize and where you can find the appropriate configuration options.

You can customize the look and feel of the following user interface elements in the admin console:

- **Sign-in pages (default and custom)**-You can customize the page that administrators see when they sign into the admin console using settings in the Authentication > Signing In > Sign-in Pages page. Using settings in this page, you can create welcome messages, sign out messages and other instructions, control page headers, customize select error messages, and create a link to a custom help page within the default system sign-in page. Or, you can upload your own custom sign-in page.
- **UI look and feel**-You can customize the header, background color, and logo displayed in the admin console using settings in the Administrators > Admin Roles > Select Role > General > UI Options page. You can also use settings in this page to enable or disable the "fly out" hierarchical menus that appear when you mouse over one of the menus in the left panel of the admin console.
- **System utilization graphs**-You can choose which system utilization graphs to display on the opening page of the admin console using settings in the System > Status > Overview page. You can also use settings in this page to fine-tune the look and data within each of the graphs.
- **Show auto-allow options**-You can show or hide the auto-allow option from yourself or other administrators who create new bookmarks for roles using settings in the Maintenance > System > Options page.
- **User role views**-You can use customization options on the Users > User Roles page to quickly view the settings that are associated with a specific role or set of roles.
- **User realm views**-You can use customization options on the Users > User Realms page to quickly view the settings that are associated with a specific user realm or set of user realms.
- **Resource policy views**-You can limit which resource policies to display on any given resource policy page based on user roles. For instance, you can configure the Users > Resource Policies > Web page of the admin console to only display those resource policies that are assigned to the "Sales" user role. You can customize these using settings in the Users > Resource Policies > Select Policy Type page of the admin console.

- **Web resource policy views**-You can limit which Web resource policy configuration pages to display using settings in Users > Resource Policies > Web > Policy Type of the admin console.
- **Administrator roles**-You can delegate select responsibilities to other administrators using settings in the Administrators > Admin Roles section of the admin console. In doing so, you can restrict the visibility of certain options and capabilities to other administrators.

Customizable End-User Interface Elements Overview

The ICS enables you to customize the look and feel of the following elements in the end-user interface:

- **Sign-in pages (default and custom)**-You can customize the page that users see when they sign into the admin console using settings in the Authentication > Signing In > Sign-in Pages page. Using settings in this page, you can create welcome messages, sign out messages and other instructions, control page headers, customize select error messages, and create a link to a custom help page within the default system sign-in page. Or, you can upload your own custom sign-in page.
- **UI look and feel**-You can customize the header, background color, and logo displayed in the admin console using settings in the Users > User Roles > Select Role > General > UI Options page. You can also use settings in this page to specify the first page the users see after they sign in, the order in which to display bookmarks, the help system to display to users, and various toolbar settings.
- **Default messages and UI look and feel**-You can specify what the default look and feel should be for all user roles using settings in Users > User Roles > [Default Options] pages of the admin console. You can also use settings in these pages to define the default errors that users see when they try to access a blocked site, SSO fails, or SSL is disabled.

REST Support for Ivanti Connect Secure

The REST API provides a standardized method for Next-Gen firewalls, NAC devices, and third-party systems to interact with ICS. Representational state transfer (REST) or RESTful Web services are one way of providing interoperability between computer systems on the Internet. REST-compliant Web services allow requesting systems to access and manipulate textual representations of Web resources using a uniform and predefined set of stateless operations. In a RESTful Web service, requests made to a resource's URI will elicit a response that may be in XML, HTML, JSON or some other defined format. ICS supports JSON format only.

REST methods determine the HTTP method for manipulating the resources defined in the service operation. The kind of operations available include those predefined by the HTTP verbs GET, POST, PUT, DELETE and so on. The response may confirm that some alteration has been made to the stored resource, and it may provide hypertext links to other related resources or collections of resources. By making use of a stateless protocol and standard operations, REST systems aim for fast performance, reliability, and the ability to grow, by re-using components that can be managed and updated without affecting the system as a whole, even while it is running.



REST API Support for ICS involves only Configuration APIs. Also, ICS supports only the GET, POST, PUT and DELETE APIs.

Authentication for REST APIs

Basic authentication using the HTTP authorization header is used to authenticate username/password on the Administrators auth. server. It is expected that the user is already configured in the Administrators auth. server. On a successful login, a random token (api_key) is generated once and sent back as a JSON response. Further access to APIs can use this api_key in their Authorization header for access.



A new random api_key is generated on a successful login. The user can continue to use this key till the administrator:

- Enables/disables the user account
- Enables/disables the Allow REST API feature for that user

The entire communication is over TLS. An example is explained below:

REQUEST

```
GET /api/v1/auth HTTP/1.1
Host: 10.209.112.106
Authorization: Basic YWRtaW5kYjpkYW5hMTIz
Content-Type: application/json
```

RESPONSE

```
HTTP/1.1 200 OK
Cache-Control: no-store
Connection: Keep-Alive
Content-Type: application/json
```

```
Expires: -1
Keep-Alive: timeout=15
{ "api_key": "p5mM1c7RQu81R2NvssLCCZhP05kf0N2ONFeYeLXX6aU=" }
```

Authorization header for all future request should perform Basic Auth using above api_key value as username and password as empty.

REQUEST

```
GET /api/v1/configuration HTTP/1.1
Host: 10.209.112.106
Authorization: Basic CDVtTWxjN1JRdTgxUjJ0dnNzTENDWmhQMDVrZjBOMk9ORmVZZUxYWDZhVT06
```

RESPONSE

```
HTTP/1.1 200 OK
Content-Length ?283
Content-Type ?application/json
{
  "administrators":
  { "href": "/api/v1/configuration/administrators" }
  ,
  "authentication":
  { "href": "/api/v1/configuration/authentication" }
  ,
  "system":
  { "href": "/api/v1/configuration/system" }
  ,
  "users":
  { "href": "/api/v1/configuration/users" }
}
```

Configuration of REST APIs

The configuration of ICS can be accessed using REST APIs. The ICS configuration is represented in a json form when accessed using REST APIs. The structure of the JSON representation is very similar to the structure of ICS XML configuration.

A new admin UI option for users under "Administrators" authserver has been added. REST API authentication would be successful only for those users who have this option enabled.

To enable this checkbox:

1. Go to **Authentication > Auth. Servers > Administrators > Update Administrator admin1**.
2. Select the **Allow access to REST APIs** checkbox. See [REST API Configuration](#)
3. Click on **Save Changes**.

REST API Configuration

Auth Servers > Administrators > Update Administrator admin1

Update Administrator admin1

Full Name:

Authenticate using: Administrators

Password:

Confirm Password:

One-time use (disable account after the next successful sign-in)

Allow console access

Allow access to REST APIs This is the new check box added

Enabled

Disabled

Quarantined

Require user to change password at next sign in

Note: You must also configure password management on the [Authentication server Settings](#) with 'Allow users inherit the server's password management capabilities.'

Enabling REST API Access for an Administrator from the Console

REST API access for an administrator user can be enabled during initial configuration and while creating a new administrator user.

During initial provisioning, there are no administrator accounts configured and the system prompts to create a new administrator user. For the option "Do you want to enable REST API access for this administrator (y/n):", enter **y**. Note that any characters other than "y" or "n" are invalid responses.

```
~ — ssh shri@10.243.53.143

Internal port configuration completed, proceeding to next step...
-----

Internal NIC: .[OK]
Currently there are no administrators configured...

Please create an administrator user.
Admin username: admindb
Password:
Confirm password:
Do you want to enable REST API access for this administrator (y/n): y

The administrator was successfully created.
-----
```

When creating a new administrator user from the console using the option "2. Create admin username and password", for the option "Do you want to enable REST API access for this administrator (y/n):", enter y.


```

~ — ssh shri@10.243.53.143

Current version: 9.0R1 (build 63950)
Rollback version: 8.3R4 (build 60528)
Reset version: 8.1R4.1 (build 37682)

Licensing Hardware ID: 0332MJ0WK0NUP111S
Serial Number: 0332122015100018

Please choose from among the following options:
 0. Start shell
100. mount root rw and start rsync...
101. mount root rw and chpax /home/bin...
102. modify platform code...
103. validate files...
104. Start sshd for debugging ...
105. Manage fault injection scenarios
 1. Network Settings and Tools
 2. Create admin username and password
 3. Display log/status
 4. System Operations
 5. Toggle password protection for the console (Off)
 6. Create a Super Admin session.
 7. System Maintenance
 8. Reset allowed encryption strength for SSL
[Choice: 2

Please create an administrator username and password.
[Admin username: consoleadmin

[Password:
[Confirm password:
[Do you want to enable REST API access for this administrator (y/n): y

The administrator consoleadmin was successfully created.

```

Sample GET/POST/PUT/DELETE Request and Responses

Below is a sample of GET/POST/PUT/DELETE request and responses:

POST API Call: Create User for Existing Local Authentication Server

REQUEST

```

POST /api/v1/configuration/authentication/auth-servers/auth-server/Sys-
Local/local/users/user HTTP/1.1
Host: 10.209.112.106
Authorization: Basic cDVtTWxjN1JRdTgxUjJOdnNzTENDWmhQMDVrZjBOMk90RmVZZUxYWDZhVT06
Content-Type: application/json

{
  "change-password-at-signin": "false",

```

```
"console-access": "false",
"enabled": "true",
"fullname": "user0001",
"one-time-use": "false",
"password-encrypted":
"3u+UR6n8AgABAAAATjgR31G4neKag2hxI+wjaNsRRZGD6wMQVv+DPQZdUrQi5IWPuihJf8tnrsBV0XCQ1
y6WgZ79Jvlfyzmssg=",
"username": "user0001"
}
```

RESPONSE

```
200 OK
Content-Length: 122
Content-Type: application/json
{
  "result": {
    "info": [
      {
        "message": "Operation succeed without warning or error!"
      }
    ]
  }
}
```

Representing Configuration Resources Using Links

When performing a GET request on a configuration resource, the json response may have "href" attributes to represent smaller resources within.

As an example, "GET /api/v1/configuration" returns:

```
{
  "users": {
    "href": "/api/v1/configuration/users"
  },
  "system": {
    "href": "/api/v1/configuration/system"
  },
  "authentication": {
    "href": "/api/v1/configuration/authentication"
  },
  "administrators": {
```

```

      "href": "/api/v1/configuration/administrators"
    }
  }
}

```

The href values can be used to access smaller resources.

GET API Call: Fetch the specific User under Local Authentication Server

REQUEST

```

GET /api/v1/configuration/authentication/auth-servers/auth-server/Sys-
Local/local/users/user/user0001 HTTP/1.1
Host: 10.209.112.106
Authorization: Basic cDVtTWxjN1JRdTgxUjJ0dnNzTENDWmhQMDVrZjBOMk90RmVZZUxYWDZhVT06
Content-Type: application/json

```

RESPONSE

```

200 OK
Content-Length: 309
Content-Type: application/json
{
  "change-password-at-signin": "false",
  "console-access": "false",
  "enabled": "true",
  "fullname": "user0001",
  "one-time-use": "false",
  "password-encrypted":
  "3u+UR6n8AgABAAAATjgR31G4neKag2hxI+wjaNsRRZGD6wMQVkJEQv+DPQZdUrQi5IWPuihJf8tnrsBV0XCQ1
y6WgZ79Jv1fyzmssg==",
  "username": "user0001"
}

```

PUT API Call: Update Fullname field of Specific user

REQUEST

```

PUT /api/v1/configuration/authentication/auth-servers/auth-server/Sys-
Local/local/users/user/user0001/fullname HTTP/1.1
Host: 10.209.112.106
Authorization: Basic cDVtTWxjN1JRdTgxUjJ0dnNzTENDWmhQMDVrZjBOMk90RmVZZUxYWDZhVT06
Content-Type: application/json
Cache-Control: no-cache

```

```
Postman-Token: 1ca1c683-4cb4-f629-53d9-cdabb9d6f092
{
  "fullname": "REST API test for user0001"
}
```

RESPONSE

```
200 OK
Content-Length: 122
Content-Type: application/json
{
  "result": {
    "info": [
      {
        "message": "Operation succeed without warning or error!"
      }
    ]
  }
}
```

After Updation fetch the User details and observe the fullname field updated:

REQUEST

```
GET /api/v1/configuration/authentication/auth-servers/auth-server/Sys-
Local/local/users/user/user0001 HTTP/1.1
Host: 10.209.112.106
Authorization: Basic cDVtTWxjN1JRdTgxUjJOdnNzTENDWmhQMDVrZjBOMk90RmVZZUxYWDZhVT06
Content-Type: application/json
RESPONSE
200 OK
Content-Length ?327
Content-Type ?application/json
{
  "change-password-at-signin": "false",
  "console-access": "false",
  "enabled": "true",
  "fullname": "REST API test for user0001",
  "one-time-use": "false",
  "password-encrypted":
"3u+UR6n8AgABAAAATjgR31G4neKag2hxI+wjaNsRRZGD6wMQVkJEQv+DPQZdUrQi5IWPuihJf8tnrsBV0XCQ1
y6WgZ79Jvlfyzmssg==",
  "username": "user0001"
```

```
}
```

DELETE API Call: DELETE Specific User

REQUEST

```
DELETE /api/v1/configuration/authentication/auth-servers/auth-server/Sys-Local/local/users/user/user0001 HTTP/1.1
Host: 10.209.112.106
Authorization: Basic cDVtTWxjN1JRdTgxUjJOdnNzTENDWmhQMDVrZjBOMk9ORmVZZUxYWDZhVT06
Content-Type: application/json
```

RESPONSE

```
200 OK
Content-Length ?122
Content-Type ?application/json
{
  "result": {
    "info": [
      {
        "message": "Operation succeed without warning or error!"
      }
    ]
  }
}
```

After deleting Try to fetch the resource and you would observe 404 response

REQUEST

```
GET /api/v1/configuration/authentication/auth-servers/auth-server/Sys-Local/local/users/user/user0001 HTTP/1.1
Host: 10.209.112.106
Authorization: Basic cDVtTWxjN1JRdTgxUjJOdnNzTENDWmhQMDVrZjBOMk9ORmVZZUxYWDZhVT06
Content-Type: application/json
Cache-Control: no-cache
Postman-Token: c94a2f29-2b52-4ed1-3987-302cbce96a30
```

RESPONSE

```
404 NOT FOUND
Content-Length: 105
```

```
Content-Type: application/json
{
  "result": {
    "errors": [
      {
        "message": "Resource does not exist."
      }
    ]
  }
}
```

FIPS Level 1 Support (Software FIPS)

Understanding Ivanti FIPS Level 1 Support

What Is FIPS?

Federal Information Processing Standard (FIPS) are a set of standards that define security requirements for products that implement cryptographic modules used to secure sensitive but unclassified information. The most recent standards are defined in the FIPS Publication 140-2.

The FIPS documents define, among other things, security levels for computer and networking equipment. U.S. Federal Government departments, and other organizations, use FIPS to evaluate the cryptographic capabilities of the equipment they consider for purchase. Cryptographic modules are validated against separate areas of the FIPS specification. An overall certification level is assigned based on the minimum level achieved in any area. Although primarily aimed at environments requiring strict security, FIPS levels are increasingly enforced as qualifying criteria for all U.S. Federal Government contracts. Security-conscious private enterprises might also use FIPS levels as an equipment evaluation benchmark. FIPS levels also serve as a customer-neutral description of vendor requirements. Vendors can engineer security products to FIPS levels and extend the applicability and eligibility of these products across a broad customer base, thereby eliminating exhaustive and time-consuming customer-by-customer product qualification procedures.

What Is FIPS Level 1 Support?

Ivanti offers FIPS level 1 support for both Ivanti Connect Secure and Ivanti Policy Secure. Both services use a 140-2 level 1 certified cryptographic module to comply with FIPS. When FIPS level 1 support is enabled applications, such as browsers, accessing the web server must support Transport Layer Security (TLS), the latest version of Secure Socket Layer (SSL). If the platform features hardware acceleration, then for SSL processing SSL hardware acceleration is disabled as hardware acceleration does not comply with FIPS validation. Only FIPS approved algorithms are used when in FIPS level 1 support is enabled.

[For more information about the Cryptography Module, see the security policy](#) and the [validation certificate](#). For a complete list of validated FIPS 140-1 and FIPS 140-2 cryptography modules, see <http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140val-all.htm#2012>.

Enabling FIPS Level 1 Support

Once you enable FIPS level 1 support, your browser is restricted to specific custom cipher strengths. A list of supported ciphers is shown during the enabling process.

When you enable FIPS level 1 support, the following events occur on the system:

- The Web server restarts and turns on FIPS level 1 support. The Web server now allows only TLSv1.0, TLSv1.1 and TLSv1.2 protocols that include FIPS approved cryptographic algorithms which include Suite B cipher suites.



Once FIPS level 1 support is enabled, new client sessions will use FIPS if the client supports FIPS. Existing client sessions may not be using FIPS. To ensure FIPS capable clients are in FIPS level 1 support, all client sessions should be terminated after the FIPS level 1 support is enabled. Administrators can use the **System > Status > Active Users** page to terminate client sessions.

- If the platform features hardware acceleration, when FIPS level 1 support is enabled SSL processing does not utilize the hardware acceleration. IPSec hardware acceleration is not affected.
- The **FIPS compliant Network Connect** option is enabled automatically on user roles. However, if this option is disabled manually for a role, and FIPS level 1 support is disabled and then reenabled, the option remains disabled for the role.

The following event logs are generated for FIPS level 1 support:

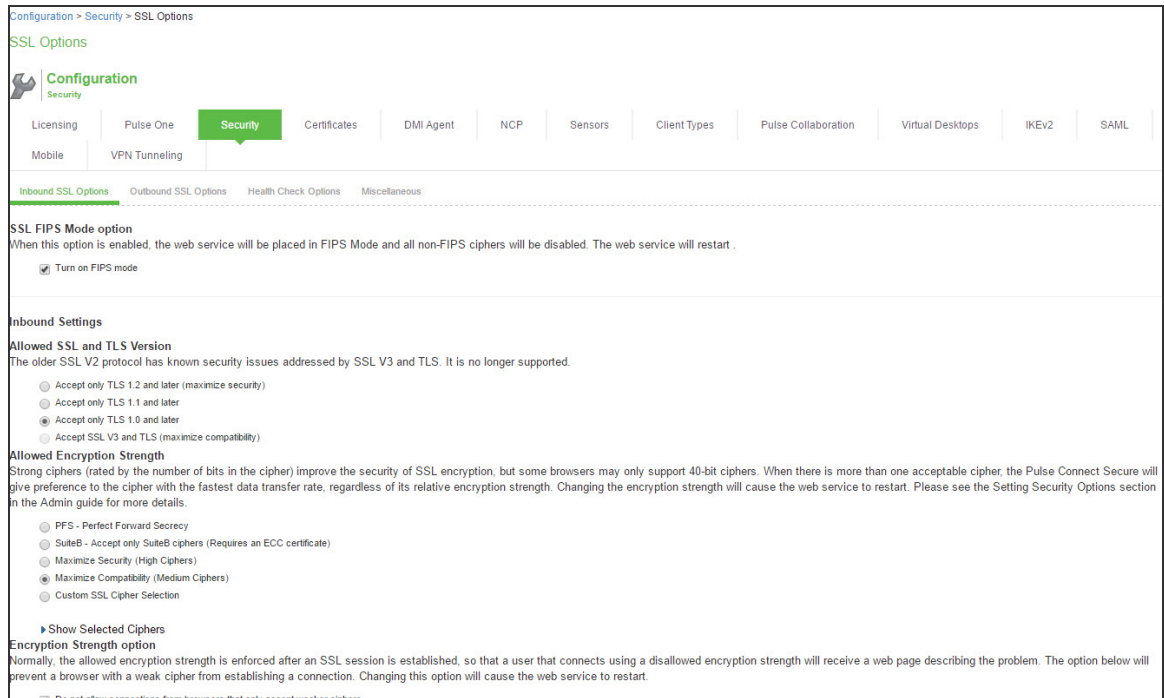
- SYS30966 when the web server turns FIPS level 1 support on.
- ADM30965 when the administrator turns FIPS level 1 support on or off.
- ERR30967 when the web server fails to turn on FIPS level 1 support.

To enable FIPS level 1 support:

1. Select **System > Configuration > Security > Inbound SSL** Options.

Under SSL FIPS Mode option, select **Turn on FIPS mode**. See [Enabling FIPS Level 1 Support](#)

Enabling FIPS Level 1 Support



Once you turn on FIPS level 1 support, the following changes are made:

- Under Allowed SSL and TLS Version, the **Accept only TLS 1.0 and later** option is selected.
 - Under Allowed Encryption Strength, the **Maximum Compatibility** Ciphers is set. See [Enabling FIPS Level 1 Support](#) Only FIPS approved algorithms are selected. All other options under this section are disabled. See [Supported Cipher Suites when FIPS Level 1 Support is Enabled and Disabled](#)
 - Under Encryption Strength, the Do not allow connections from browsers that only accept weaker ciphers option is selected. You cannot disable this selection.
2. Click **Save Changes**.
 3. Entries are made in the Events logs (see [Events Log Entries for FIPS Level 1 \(System > Log/Monitoring > Events > Logs\)](#)) and Admin Access logs (see [Admin Access Logs for FIPS Level 1 Encryption Strength Changes \(Main Menu: System\)](#)) to show that FIPS level 1 support is enabled.

Events Log Entries for FIPS Level 1 (**System > Log/Monitoring > Events > Logs**)

Turning Off FIPS Level 1 and Resetting Encryption Strength from the Serial Console

Please choose the operation to perform:

1. Network Settings and Tools
2. Create admin username and password
3. Display log/status
4. System Operations
5. Toggle password protection for the console (Off)
6. Create a Super Admin session.
7. System Maintenance
8. Turn off FIPS Mode and reset allowed encryption strength for SSL
Choice: 8



Once you turn off FIPS level 1 support, option 8 is relabeled "Reset allowed encryption strength for SSL."

Installing a Self-Signed Certificate from the Serial Console

<p>Problem</p>	<p>Description: An administrator can be locked out of the system if their browser does not support the certificate assigned to the network port. For example, if your system has an ECC certificate assigned to the internal port and your browser does not support ECC certificates you cannot log in to the device using the internal port.</p>
<p>Solution</p>	<p>You can use the serial console to create and install a self-signed RSA certificate onto the internal port to allow access. Once you connect to the serial console, select option 4. System Operations followed by Option 7. Install self-signed certificate. It may take a few minutes for the 2048-bit key size self-signed certificate to be created and installed on your device. Once the certificate is installed, you can now log in to the device.</p>

Creating and Installing an RSA Certificate from the Serial Console

Please choose the operation to perform:

1. Network Settings and Tools
2. Create admin username and password
3. Display log/status
4. System Operations
5. Toggle password protection for the console (Off)
6. Create a Super Admin session.
7. System Maintenance
8. Turn off FIPS Mode and reset allowed encryption strength for SSL
Choice: 4

Please choose the operation to perform:

1. Reboot this Ivanti Connect Secure
2. Shutdown this Ivanti Connect Secure
3. Restart services at this Ivanti Connect Secure
4. Rollback this Ivanti Connect Secure
5. Factory reset this Ivanti Connect Secure
6. Clear all configuration data at this Ivanti Connect Secure
7. Install self-signed certificate

Choice: 7

Are you sure you want to install a newly-created RSA self-signed certificate on the internal port? (y/n) y

Please provide information to create a self-signed Web server

digital certificate.

Common name (example: secure.company.com) :
myname.mycompany.com

Organization name (example: Company Inc.): MyCompany Inc.

Please enter some random characters to augment the system's

random key generator. We recommend that you enter approximately

thirty characters.

Random text (hit enter when done):abcdef1234567

Creating self-signed digital certificate - this may take several minutes...

The self-signed digital certificate was successfully created.

Supported Cipher Suites when FIPS Level 1 Support is Enabled and Disabled

The tables in this topic list the cipher suites that are supported by the web server when the FIPS level 1 support is disabled and enabled.

Supported Cipher Suites When FIPS Level 1 Support is Disabled

When the FIPS level 1 support is disabled and when SSL hardware acceleration is not present or is disabled, the web server gives preference to cipher suites that use RC4 for bulk encryption. When SSL hardware acceleration is present and enabled, the web server gives preference to AES128, AES256 and 3DES over RC4, in that order.

The following table lists the Supported Cipher Suites With FIPS Level 1 Support Off, Hardware Acceleration Enabled and RSA Server Certificates In Use:

Cipher Suite	Protocol
TLS_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_RSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_RSA_WITH_AES_256_CBC_SHA256	TLSv1.2
TLS_RSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_RSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later

Cipher Suite	Protocol
TLS_RSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_RSA_WITH_RC4_128_MD5	SSLv3 and later

The following table lists the Supported Cipher Suites with FIPS Level 1 Support Off, Hardware Acceleration Enabled and ECC Server Certificates in Use:

Cipher Suite	Protocol
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2

Cipher Suite	Protocol
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDHE_ECDSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_RC4_128_SHA	SSLv3 and later

The following table lists the Supported Cipher Suites with FIPS Level 1 Support Off, Hardware Acceleration Disabled and RSA Server Certificates in Use:

Cipher Suite	Protocol
TLS_RSA_WITH_RC4_128_SHA	SSLv3 and later

Cipher Suite	Protocol
TLS_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_RSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_RSA_WITH_AES_256_CBC_SHA256	TLSv1.2
TLS_RSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_RSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDHE_RSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_RSA_WITH_RC4_128_MD5	SSLv3 and later

The following table lists the Supported Cipher Suites with FIPS Level 1 Support Off, Hardware Acceleration Disabled and ECC Server Certificates in Use:

Cipher Suite	Protocol
TLS_ECDHE_ECDSA_WITH_RC4_128_SHA	SSLv3 and later

Cipher Suite	Protocol
TLS_ECDH_RSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384	TLSv1.2

Cipher Suite	Protocol
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later

Supported Cipher Suites When FIPS Level 1 Support is Enabled

When FIPS level 1 support is enabled, only TLSv1.0, v1.1, v1.2 and AES256, 3DES and AES128 are allowed. The order of the cipher suites is not dependent on the SSL hardware acceleration module since hardware acceleration is not used when FIPS level 1 support is enabled.

When FIPS level 1 support is enabled, the following settings are automatically configured:

- In the SSL Options window:
 - Under Allowed SSL and TLS Version, the **Accept only TLS 1.0** and later option is selected. All other options under this section are disabled.
 - Under Allowed Encryption Strength, the **Maximum Compatibility** ciphers option is selected. Only FIPS approved algorithms are selected. All other options under this section are disabled.
 - Under Encryption Strength Option, the **Do not allow connections from browsers that only accept weaker ciphers option is selected.**

- **SSL hardware acceleration** is disabled. IPsec hardware acceleration is not affected by the FIPS level 1 support being enabled.

In the following table, the first four cipher suites are given preference due to the requirements in RFC 6460. The first two cipher suites meeting the requirement for Suite B Profile for TLS 1.2. The next two meeting the requirement for Suite B Transitional Profile for TLS 1.0 and 1.1.

The following table lists the Supported Cipher Suites with FIPS Level 1 Support on and ECC Server Certificates in Use:

Cipher Suite	Protocol
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA	TLSv1.0 and later
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA	TLSv1.0 and later
TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA	TLSv1.0 and later
TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2

Cipher Suite	Protocol
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA	TLSv1.0 and later
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA	TLSv1.0 and later

The following table lists the Supported Cipher Suites with FIPS Level 1 Support on and RSA Server Certificates in Use:

Cipher Suite	Protocol
TLS_RSA_WITH_AES_256_CBC_SHA256	TLSv1.2
TLS_RSA_WITH_AES_256_CBC_SHA	TLSv1.0 and later
TLS_RSA_WITH_3DES_EDE_CBC_SHA	TLSv1.0 and later
TLS_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_RSA_WITH_AES_128_CBC_SHA	TLSv1.0 and later
TLS_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2

Compression

About Compression

The system improves performance by compressing common types of Web and file data such as HTML files, Word documents, and images.

The system determines whether it should compress the data accessed by users by using the following process:

1. The system verifies that the accessed data is a compressible type. Compressing many common data types such as HTML files, and Word documents is supported.
2. If the user is accessing Web data, the system verifies that the user's browser supports compression of the selected data type.
The system determines compression supportability based on the browser's user-agent and the accept-encoding header. It supports the compression of all of the standard Web data types if it determines that the user-agent is compatible with Mozilla 5, Internet Explorer 5, or Internet Explorer 6. The system supports only compressing HTML data, however, if it determines that the browser's user-agent is only compatible with Mozilla 4.
3. The system verifies that compression is enabled at the system level. You can enable system-level compression through the Maintenance > System > Options page of the admin console.
4. The system verifies that compression resource policies or autopolicies are enabled for the selected data type and comes with resource policies that compress data. You may enable these policies or create your own through the following pages of the admin console:
 - Users > Resource Policies > Web > Compression.
 - Users > Resource Policies > Files > Compression.

You may also create resource profile compression autopolicies through the Users > Resource Profiles > Web > Web Applications/Pages page of the admin console.

If all of these conditions are met, the system runs the appropriate resource policy either compresses or does not compress the data accessed by the user based on the configured action.

If all of these conditions are not met, the system does not run the appropriate resource policy and no resource policy items appear in the log files.

The system comes pre-equipped with three resource policies that compress Web and file data. If you are upgrading from a pre-4.2 version of the system software and you previously had compression enabled, these policies are enabled. Otherwise, if you previously had compression disabled, these policies are disabled.

The Web and file resource policies created during the upgrade process specify that the system should compress all supported types of Web and File data, including types that were not compressed by previous versions of the appliance. All data types that were not compressed by previous product versions are marked with an asterisk (*) in the supported data types list below.

The system supports compressing the following types of Web and file data:

- text/plain (.txt)
- text/ascii (.txt)*
- text/html (.html, .htm)
- text/css (.css)
- text/rtf (.rtf)
- text/javascript (.js)
- text/xml (.xml)*
- application/x-javascript (.js)
- application/msword (.doc)
- application/ms-word (.doc)*
- application/vnd.ms-word (.doc)*
- application/msexcel (.xls)*
- application/ms-excel (.xls)*
- application/x-excel (.xls)*
- application/vnd.ms-excel (.xls)*
- application/ms-powerpoint (.ppt)*
- application/vnd.ms-powerpoint (.ppt)*



The data types denoted by an asterisk * were not compressed by pre-4.2 versions of the system software.

Also note that the system does not compress files that you upload-only files that you download from the system.

Additionally, the system supports compressing the following types of files:

- text/html (.html, .htm)
- application/x-javascript (.js)
- text/javascript (.js)
- text/css (.css)
- application/perl (.cgi)

Enabling System-Level Compression

To enable system-level compression:

1. Select **Maintenance > System > Options**.
2. Select the **Enable gzip compression** check box to reduce the amount of data sent to browsers that support HTTP compression. Note that after you enable this option, you must also configure Web and file resource policies specifying which types of data the system should compress.
3. Click **Save Changes**.

Localization

About Multi-Language Support for Ivanti Connect Secure

The system provides multi-language support for file encoding, end-user interface display, and customized sign-in and system pages. It supports the following languages:

- English (US)
- Chinese (Simplified)
- Chinese (Traditional)
- French
- German
- Japanese
- Korean
- Spanish
- Italian
- Polish

Encoding Files for Multi-Language Support

Character encoding is a mapping of characters and symbols used in written language into a binary format used by computers. Character encoding affects how you store and transmit data. The encoding option in **Users > Resource Policies > Files > Encoding** allows you to specify the encoding to use when communicating with Windows and NFS file shares. The encoding option does not affect the end-user language environment.

To specify the internationalization encoding for system traffic:

1. In the admin console, choose **Users > Resource Policies > Files > Encoding**.
2. Select the appropriate option:
 - Western European (ISO-8859-1) (default) (Includes English, French, German, Spanish)
 - Simplified Chinese (CP936)

- Simplified Chinese (GB2312)
- Traditional Chinese (CP950)
- Traditional Chinese (Big5)
- Japanese (Shift-JIS)
- Korean

3. Click **Save Changes**.

Localizing the User Interface

The system provides a means to display the end-user interface in one of the supported languages. Combining this feature with (custom) sign-in and system pages and a localized operating system provides a fully localized user experience.

When you specify a language, the system displays the user interface, including all menu items, dialogs generated by the system, and the help file in the chosen language for all users regardless of which realm they sign in to.

To configure localization options:

1. In the admin console, choose **Maintenance > System > Options**.
2. Use the End-user Localization drop-down list to specify the language in which to display the end-user interface (optional). If you do not specify a language, the end-user interface displays based on the settings of the browser.
3. Click **Save Changes**.

Localizing Custom Sign-In and System Pages

The system provides several zip files that contain different sets of sample template files for various pages that may appear during the sign-in process. Use these template files along with the template toolkit language to create localized custom sign-in and system pages for your end users.

Editing the default sign-in page using text in the language of your choice is a quick way to provide your users with a localized sign-in page.

Smart Phones

Smart Phones

In addition to allowing users to access the system from standard workstations and kiosks, the system also allows end users access from connected PDAs, handhelds and smart phones such as i-mode and Pocket PC. When a user connects from a PDA or handheld device, the system determines which pages and functionality to display based on settings in the System > Configuration > Client Types page of the admin console. By default, settings in this page specify that when accessing the system using a(n):

- **i-mode device**-The system displays compact HTML (cHTML) pages without tables, images, JavaScript, Java, or frames to the user. Depending on which features you enable through the admin console, the end user may browse the Web, link to Web bookmarks, single sign-on to other applications, and edit their preferences (including clearing their cache and editing their system/LDAP password). The system allows i-mode users to access supported features using access keys on their phone's keypad as well as through standard browse-and-select navigation.
- **Pocket PC device**-The system displays mobile HTML pages with tables, images, JavaScript and frames, but does not process Java. Depending on which features you enable through the admin console, the end user may access Mobile Notes and OWA e-mail applications, browse the Web, link to Web bookmarks, single sign-on to other applications, and edit their preferences (including clearing their cache and editing their system/LDAP password).

PDA and handheld users cannot access the admin console or most of the system's advanced options, including file browsing, VPN Tunneling, Host Checker, since PDA and handheld devices do not generally support the ActiveX, Java, or JavaScript controls on which these features depend.

Also note that i-mode users cannot access cookie-based options, including session cookies, since most i-mode browsers do not support HTTP cookies. The system rewrites hyperlinks to include the session ID in the URL instead of using cookies. The system reads the session ID when the user accesses the URL.



In order to improve the response time, the following icons are not displayed when accessing the home page: help, sign out, open bookmark in new page, and PSAM.

Configuring Connect Secure for PDAs and Handhelds

To properly configure the system to work with PDAs and handheld devices, you must:

1. **Enable access at the system level**-If you want to support browsers other than the defaults provided with the system, you must enter the user agent strings of the PDA and handheld operating systems that you want to support in the System > Configuration > Client Types tab. For a complete list of supported PDA and handheld browsers, see the Supported Platforms document posted on the Support web site.
2. **Evaluate your user roles and resource policies**-Depending on which Ivanti Connect Secure features you have enabled, you may need to either modify your existing roles and resource policies for PDA and handheld users or create new ones. Note that:
 - Mobile device users cannot access roles or policies that require Host Checker since handheld devices do not generally support the ActiveX, Java, or JavaScript controls on which these features depend. You can disable these options through the following tabs:
 - Users > User Roles > *Role* > General > Restrictions
 - Resource Policies > Web > Access > Web ACL > *Policy* > Detailed Rules
 - Mobile device users may have trouble reading long role names on their small screens. If you require users to pick from a list of roles when they sign in, you may want to shorten role names in the Users > User Roles > Role > General > Overview tab.
 - Mobile device users may have trouble reading long bookmark names on their small screens. You can edit Web bookmarks in the following tabs:
 - Users > Resource Profiles > Web Application Resource Profiles > *Profile* > Bookmarks
 - Users > User Roles > *Role* > Web > Bookmarks
 - Resource Policies > Web > Access > Web ACL > *Policy* > General
 - Although advanced features such as file browsing are not supported for PDAs and handhelds, you do not need to disable them in the roles and resource policies used by mobile device users. The system simply does not display these options to mobile device users.
3. **Evaluate your authentication and authorization servers**-The system supports all of the same authentication and authorization servers for PDA and handheld users as standard users.
4. **Evaluate your realms**-Depending on which system features you have enabled, you may need to either modify your existing realms for PDA and handheld users or create new ones. Note that:

- Mobile device users cannot access the system when they try to sign into a realm that requires Host Checker since handheld devices do not generally support the ActiveX, Java, or JavaScript controls on which these features depend. You can disable these options through sub-tabs in the System > Configuration > Security page.
 - Mobile device users may have trouble reading long realm names on their small screens. If you require users to pick from a list of realms when they sign in, you may want to shorten realm names in the Users > User Realms > *Realm* > General tab.
5. **Evaluate your sign-in policy to use**-If you want to use a different sign-in page for Pocket PC users, you can define it in the Authentication > Signing In > Sign-in Pages tab and then create a sign-in policy that references the page using options in the Authentication > Signing In > Sign-in Policies tab. Or, you can create a custom sign-in page using the Pocket PC template files that are available in sample.zip.
 6. **Specify allowed encryption strength**-Different types of devices allow different encryption strengths. You should specify the encryption strength in Ivanti Connect Secure to match the requirement of your devices. For example, mobile phones often only accept 40-bit encryption. Review your end-users' device requirements and specify the allowed encryption strength on the System > Configuration > Security tab.

Defining Client Types

The Client Types tab allows you to specify the types of systems your users may sign in from and the type of HTML pages to display when they do. In addition, client types are used to identify the operating system shown on the Device Management page for devices that use ActiveSync to synchronize e-mail with a Microsoft Exchange server. The user agent string used to identify a device during login may be different from the one in the ActiveSync message. For example, in the list of default user agent strings, *Apple-iPhone* and *Apple-iPad* are used only in ActiveSync messages.

To manage the client types:

1. In the admin console, choose **System > Configuration > Client Types**.

2. In the **User-agent string pattern** text box, enter the user agent string for the operating system (s) that you want to support. You can specify all or part of the string. For example, you can use the default `*DoCoMo*` string to apply to all DoCoMo operating systems, or you can create a string such as `DoCoMo/1.0/P502i/c10` to apply to a single type of DoCoMo operating system. You can use the `*` and `?` wildcard characters in the string. Note that user agent strings on the system are case-insensitive.

If a device operating system shown on the Device Management page is Other, the ActiveSync message for the device has a user-agent string that is not defined here. To add the missing user-agent string:

3. Select **System > Log/Monitoring > User Access > Log**.
4. Search the User Access Log using the filter `id='AUT31094' && user='username'`. The AUT31094 is the ActiveSync log message ID, and you can select **System > Status > Devices** to get the device's username from the Device Management page. The log message looks like the following:

Device record created for user jsmith@asglab.onmicrosoft.com to obtain Authorization Only access. (activesync_id=SAMSUNG1355815045478007_AM, user-agent=SAMSUNG-SAMSUNG-SGH-I997/100.202)

1. Copy the `user-agent=` value from the log message to the User-agent string pattern text box.
 2. Select the client type (see Step 3) and click Add.
5. Select the type of HTML to display to users who sign in from the operating system specified in the previous step. Options include:
 - **Standard HTML**-The system displays all standard HTML functions, including tables, full-size graphics, ActiveX components, JavaScript, Java, frames, and cookies. Ideal for standard browsers, such as Firefox, Mozilla, and Internet Explorer.
 - **Compact HTML (iMode)**-The system displays small-screen HTML-compatible pages. This mode does not support cookies or the rendering of tables, graphics, ActiveX components, JavaScript, Java, VB script, or frames. (The only difference between this option and the Smart Phone HTML Basic option is the user interface.) Ideal for iMode browsers.



Form Post SSO is not supported on iMode appliances.

- **Mobile HTML (Pocket PC)**-The system displays small-screen HTML-compatible pages that may contain tables, small graphics, JavaScript, frames, and cookies, but this mode does not facilitate the rendering of java applets or ActiveX components. Ideal for Pocket PC browsers.
- **Smart Phone HTML Advanced**-The system displays small-screen HTML-compatible pages that may contain tables, small graphics, frames, cookies, and some JavaScript, but this mode does not facilitate the rendering of java applets, ActiveX components, or VB scripts. Ideal for Treo and Blazer browsers.
- **Smart Phone HTML Basic**-The system displays small-screen HTML-compatible pages. This mode does not support cookies or the rendering of tables, graphics, ActiveX components, JavaScript, Java, VB script, or frames. (The only difference between this option and the Compact HTML option is the user interface.) Ideal for Opera browsers on Symbian.



The system rewrites hyperlinks to include the session ID in the URL instead of using cookies.

- **Mobile Safari, Android, Symbian, iPad**-The Mobile Safari (iPhone/iPod Touch), Android, and Symbian selections have Basic, Advanced, and Full HTML options.
6. Specify the order that you want to evaluate the user agents. The system applies the first rule in the list that matches the user's system. For example, you may create the following user agent string/HTML type mappings in the following order
- User Agent String: *DoCoMo* Maps to: Compact HTML
 - User Agent String: DoCoMo/1.0/P502i/c10 Maps to: Mobile HTML

If a user signs in from the operating system specified in the second line, the system will display compact HTML pages to him, not the more robust mobile HTML, since his user agent string matches the first item in the list.

To order mappings in the list, select the check box next to an item and then use the up and down arrows to move it to the correct place in the list.

7. Select the Enable password masking for Compact HTML check box if you want to mask passwords entered in iMode and other devices that use compact HTML. (Devices that do not use compact HTML mask passwords regardless of whether or not you select this check box.) Note that if your iMode users' passwords contain non-numeric characters, you must disable password masking because iMode devices only allow numeric data in standard password fields. If you disable masking, passwords are still transmitted securely, but are not concealed on the user's display.

8. Click **Save Changes**.

Enabling ActiveSync for Handheld Devices

Using ActiveSync, you can synchronize data between a Windows-based desktop computer and handheld devices. Connect Secure can be used as a reverse proxy to allow users to synchronize their data without installing an additional client application on their handheld devices. More than 1000 concurrent connections are supported on a PSA7000.

Please note the following:

- Supports Windows Phone 5.0, 6.0 and 8.0 only.
- Supports Exchange Server 2010, 2013.
- ActiveSync does not use up concurrent user licenses, even when configured with certificate authentication.
- Both NTLM & Basic Auth on the Exchange server are supported.
- Both HTTP and HTTPS between Connect Secure and an Exchange server are supported.
- If Connect Secure is used for OWA & ActiveSync, the hostnames for OWA access and ActiveSync must be different.
- Direct Push is supported with ActiveSync, however you must set HTTPServerTimeout to 20 minutes or less. Direct Push is a feature built into Exchange Server 2010.
- ActiveSync does not work through a back-end web proxy.
- VIP sourcing settings are ignored for ActiveSync sessions. ActiveSync traffic from Connect Secure to a backend server is always sent with the Internal Port's source IP address.

To configure the system as a reverse proxy for use with ActiveSync:

1. In the admin console, choose Authentication > Signing In > Sign-in Policies.
2. To create a new authorization only access policy, click New URL and select authorization only access. Or, to edit an existing policy, click a URL in the Virtual Hostname column.

3. In the Virtual Hostname field, enter the name that maps to the system IP address. The name must be unique among all virtual hostnames used in pass-through proxy's hostname mode. The hostname is used to access the Exchange server entered in the Backend URL field. Do not include the protocol (for example, http:) in this field.

For example, if the virtual hostname is myapp.ivehostname.com, and the backend URL is http://www.xyz.com:8080/, a request to https://myapp.ivehostname.com/test1 by the system is converted to a request to http://www.xyz.com:8080/test1. The response of the converted request is sent to the original requesting web browser.

4. In the **Backend URL** field, enter the URL for the Exchange server. You must specify the protocol, hostname and port of the server. For example, <http://www.mydomain.com:8080/>.

When requests match the hostname in the Virtual Hostname field, the request is transformed to the URL specified in the Backend URL field. The client is directed to the backend URL unaware of the redirect.

5. Enter a Description for this policy (optional).
6. Select the server name or No Authorization from the Authorization Server drop-down menu. If you select a server, ensure that the front-end server provides the SMSESSION cookie otherwise you will receive an error.
7. Select a user role from the Role Option drop-down menu.

Only the following user role options are applicable for *Autosync*.

- HTTP Connection Timeout (Users > User Roles > *RoleName* > Web > Options > View advanced options)
- Allow browsing untrusted SSL web sites (Users > User Roles > *RoleName* > Web > Options > View advanced options)
- Source IP restrictions (Users > User Roles > *RoleName* > General > Restrictions)
- Browser restrictions (Users > User Roles > *RoleName* > General > Restrictions)

Ensure the user role you select has an associated Web Access policy.

8. Select the **Allow ActiveSync Traffic only** option to perform a basic of validation of the HTTP header to ensure the request is consistent with ActiveSync protocol. If you select this option only ActiveSync protocol requests can be processed. If validation fails, a message is created in the user's event log. If you do not select this option, both ActiveSync and non-ActiveSync requests are processed.

9. Click **Save Changes**.

The System Status Overview page displays the number of current active concurrent connections and a histogram of the active concurrent connections (Authorization Only Access Active Connections plot in the Concurrent SSL Connections graph).

Custom Expressions and System Variables

Using Custom Expressions in Rule Configuration

This topic describes custom expressions. It is intended for advanced users.

Custom Expressions

Many system rules, such as role mapping rules or resource policy rules, support custom expressions. A custom expression is a combination of variables that the system evaluates as a Boolean object. The expression returns true, false, or error.

You can write custom expressions in the following formats. Note that elements of these formats are described in greater detail in the table that follows:

- *variable comparisonOperator variable*
- *variable comparisonOperator simpleValue*
- *variable comparisonOperator (simpleValue)*
- *variable comparisonOperator (OR Values)*
- *variable comparisonOperator (AND Values)*
- *variable comparisonOperator (time TO time)*
- *variable comparisonOperator (day TO day)*
- *isEmpty (variable)*
- *isUnknown (variable)*
- *(customExpr)*
- *NOT customExpr*
- *! customExpr*
- *customExpr OR customExpr*
- *customExpr || customExpr*
- *customExpr AND customExpr*

- *customExpr && customExpr*

 The custom expression should be less than 64K.

Custom Expression Elements

The following table describes the Custom Expression Elements:

Element	Description
variable	<p>Represents a system variable. A variable name is a dot-separated string, and each component can contain characters from the set [a-z, A-Z, 0-9_] but cannot start with a digit [0-9]. Variable names are case-insensitive. For system variables that you may use in role mapping rules and resource policies.</p> <p>When writing a custom expression in a log query field, you need to use system log variables. These variables are described in the Filter Variables Dictionary on the Filter page (System > Log/Monitoring > Events User Access Admin Access > Filters > Select Filter tab).</p>
	<p>Quoting syntax for variables:</p> <p>The system supports a quoting syntax for custom expression variables that allows you to use any character except '.' (period) in a user attribute name. To escape characters in an attribute name, quote some or all of the variable name using {} (curly-braces). For example, these expressions are equivalent:</p> <pre>userAttr.{Login-Name} = 'xyz' userAttr.Login{-}Name = 'xyz' {userAttr.Login-Name} = 'xyz' userA{ttr.L}{ogin-}Name = 'xyz'</pre>
	<p>Escape characters supported within quotes:</p> <pre>\\-Escape a backslash (\). \{-Escape a left curly brace ({}). \}-Escape a right curly brace (}). \hh-Escape a hexadecimal value where hh is two characters from [0-9A-Fa-f].</pre>
	<p>Examples:</p> <pre>userAttr.{Tree Frog} = 'kermit' userAttr.{Tree\20Frog} = 'kermit'</pre>

Element	Description
	<p>There is no limit to the number of quotes you can use in a variable name. You can use the quoting syntax with any variable, not just userAttr.* variables.</p> <p>You need to use curly-brace quotes only when writing custom expressions. Usage of period works for user attributes role mapping only but may fail to pass value to the backend server. The period is mainly used as an internal separator . For example: userAttr.http://schemas.microsoft.com/ws/2008/06/identity/claims/group s</p>
<i>comparisonOperator</i>	<p>One of the following:</p> <ul style="list-style-type: none"> =-Equal to. Use with strings, numbers, and DNs. !=-Not equal to. Use with strings, numbers, and DNs. <-Less than. Use with numbers. <=-Less than or equal to. Use with numbers. >-Greater than. Use with numbers. >=-Greater than or equal to. Use with numbers.
<i>simpleValue</i>	<p>One of the following:</p> <ul style="list-style-type: none"> string - quoted string that may contain wildcards. IP Address-a.b.c.d subnet-a.b.c.d/subnetBitCount or a.b.c.d/netmask number-Positive or negative integer day-SUN MON TUE WED THU FRI SAT <p>Notes about strings:</p> <ul style="list-style-type: none"> A string may contain all characters except <nl> (newline) and <cr> (carriage return). Strings can be any length. String comparisons are case-insensitive. Strings can be quoted with single- or double-quotes. A quoted string may contain wildcards, including star(*), question mark (?), and square brackets ([]). variable comparisonOperator variable comparisons are evaluated without wildcard matching. Use a backslash to escape these characters: single-quote (') - \' double-quote (") - \" backslash (\) - \\ hexadecimal - \hh [0-9a-fA-F] <p>Note about day:</p>

Element	Description
	<p>Day and time comparisons are evaluated in the system's time zone. Day range (day TO day) calculations start with the first day and step forward until the second day is reached. In time range (time TO time) calculations, the first value must be earlier than the second value. Only time variables can be compared to day and time values. The time variables are: <code>time.*</code> and <code>loginTime.*</code>.</p>
<i>time</i>	<p>Time of day in one of the following formats: HH:MM - 24-hour HH:MMam - 12-hour HH:MMpm - 12-hour H:MM - 24-hour H:MMam - 12-hour H:MMpm - 12-hour</p> <p>Day and time comparisons are evaluated in the system's time zone. Day range (day TO day) calculations start with the first day and step forward until the second day is reached. In time range (time TO time) calculations, the first value must be earlier than the second value. Only time variables can be compared to day and time values. The time variables are: <code>time.*</code> and <code>loginTime.*</code>.</p>
<i>OR Value</i>	<p>String containing one or more OR comparisons: Examples: <code>variable comparisonOperator (number OR number ...)</code> <code>variable comparisonOperator (string OR string ...)</code></p>
<i>AND Value</i>	<p>String containing one or more AND comparisons. Examples: <code>variable comparisonOperator (number AND number ...)</code> <code>variable comparisonOperator (string AND string ...)</code></p>
<i>isEmpty</i>	<p>Function that takes a single variable name (variable) argument and returns a boolean value. <code>isEmpty()</code> is true if the variable is unknown or has a zero-length value, zero-length strings, and empty lists. Example: <code>isEmpty(userAttr.terminationDate)</code></p>
<i>isUnknown</i>	<p>Function that takes a single variable name (variable) argument and returns a boolean value. <code>isUnknown()</code> is true if the variable is not defined. User attributes (<code>userAttr.*</code> variables) are unknown if the attribute is not defined in LDAP or if the attribute lookup failed (such as if the LDAP server is down).</p>

Element	Description
	Example: <code>isUnknown(userAttr.bonusProgram)</code>
<i>NOT, !</i>	Logical negation comparisonOperator. The negated expression evaluates to true if the customExpr is false and evaluates to false if the customExpr is true. The operators NOT, AND, and OR are evaluated from highest to lowest precedence in this order: NOT (from right), AND (from left), OR (from left).
<i>OR, </i>	Logical operator OR or , which are equivalent. The operators NOT, AND, and OR are evaluated from highest to lowest precedence in this order: NOT (from right), AND (from left), OR (from left).
<i>AND, &&</i>	Logical AND or &&, which are equivalent. The operators NOT, AND, and OR are evaluated from highest to lowest precedence in this order: NOT (from right), AND (from left), OR (from left).
<i>customExpr</i>	Expression written in the Custom Expression Syntax (see above).

Wildcard Matching

In a quoted string, supported wildcards include:

- star (*)-A star matches any sequence of zero or more characters.
- question mark (?)-A question mark matches any single character.
- square brackets ([])-Square brackets match one character from a range of possible characters specified between the brackets. Two characters separated by a dash (-) match the two characters in the specified range and the lexically intervening characters. For example, 'dept[0-9]' matches strings "dept0", "dept1", and up to "dept9".

To escape wildcard characters, place them inside square brackets. For example, the expression 'userAttr.x = " value [*]" ' evaluates to true if attribute x is exactly "value*".

Using Multivalued Attributes

Multivalued attributes-attributes that contain two or more values-provide you with a convenient method for defining resources that expand into multiple individual bookmarks on the users' bookmarks page.

For example, assume that the user's LDAP directory contains the multivalued attribute HomeShares: \\Srv1\Sales;\\Srv2\Marketing. When you configure the Windows File share resource definition using the HomeShares multivalued attribute, \\<userAttr.HomeShares>, the user sees two bookmarks:

- \\Srv1\Sales
- \\Srv2\Marketing

Now let's assume the user's LDAP directory contains a second multivalued attribute defined as HomeFolders: Folder1;Folder2;Folder3. When you configure the Windows File share resource using both of the multivalued attributes, \\<userAttr.HomeShares>\<userAttr.HomeFolders>, the user sees the following six bookmarks:

- \\Srv1\Sales\Folder1
- \\Srv1\Sales\Folder2
- \\Srv1\Sales\Folder3
- \\Srv2\Marketing\Folder1
- \\Srv2\Marketing\Folder2
- \\Srv2\Marketing\Folder3

The only exception to this functionality is when the variable includes an explicit separator string. In this case, only one bookmark containing multiple resources displays on the users' bookmark page.

You specify the separator string in the variable definition using the syntax `sep='string'` where string equals the separator you want to use. For example, to specify a semi-colon as the separator, use the syntax `<variable.Attr sep='; '>`.

Use the following syntax for multivalued attributes handling. Note that `<variable>` refers to a session variable such as `<userAttr.name>` or `<CertAttr.name>`:

- `<variable[Index]>` - You specify indexes in a variety of ways. If, for example, the total number of values for a given index is 5, and you want to specify the entire range of values you use `<variable[ALL]>`. If you want to specify only the fourth value, you use `<variable[4]>`.
- `<variable>` is the same as `<variable[ALL]>`.
- `<variable>` is the same as `<variable[ALL]>`.
- `<variable sep='str'>` and `<variable[All] sep='str'>` - These variable definitions always refer to a single string value with all the tokens expanded out with separator strings between the values.



Variable names cannot contain spaces.

Specifying Multivalued Attributes in a Bookmark Name

Another common case of using multivalued attributes occurs when you include a variable in a bookmark name and in a URL or file server/share field.

For example, again assume that the user's LDAP directory contains the multivalued attribute HomeShares: \\Srv1\Sales;\\Srv2\Marketing. When you configure the Windows File share resource definition using the HomeShares multivalued attribute, \\<userAttr.HomeShares>, and you use the same attribute in the bookmark name field, <userAttr.HomeShares>, the system creates two bookmarks:

- Srv1\Sales bookmark pointing to \\Srv1\Sales
- Srv2\Marketing bookmark pointing to \\Srv2\Marketing

This does not create a situation in which you end up with the following set of conditions:

- Srv1\Sales bookmark pointing to \\Srv1\Sales
- Srv1\Marketing bookmark pointing to \\Srv1\Marketing (error)
- Srv2\Sales bookmark pointing to \\Srv1\Sales (error)
- Srv2\Marketing bookmark pointing to \\Srv2\Marketing

Distinguished Name Variables

You can compare a distinguished name (DN) to another DN or to a string, but the system ignores wildcards, white space, and case. Note, however, that the system takes the order of DN keys into consideration.

When the system compares an expression to a DN to a string, it converts the string to a distinguished name before evaluating the expression. If the system cannot convert the string due to bad syntax, the comparison fails. The DN variables are:

- userDN
- certDN
- certIssuerDN

The system also supports DN suffix comparisons using the **matchDNSuffix** function. For example:


```
matchDNSuffix( certDn, "dc=danastreet,dc=net")
```

Within the parenthesis, the first parameter is the " full" DN and the second is the suffix DN. You can use a variable or string for each parameter. Note that this first parameter should have more keys than the second (suffix parameter). Otherwise, if they are equal, it is the same as <firstparam> = <secondparam>. If the second parameter has more keys, **matchDNSuffix** returns false.

System Variables

The following table lists and defines system variables, gives an example for each system variable, and provides a guide as to where you may use system variables.

The following table lists the System Variables and Examples:

Variable	Description		Examples
authMethod	Type of authentication method used to authenticates a user.	role mapping rules, resource policy rules	authMethod = 'ACE Server'
certAttr.<cert-attr>	Attributes from a client-side certificate. Examples of certAttr attributes include: C - country CN - common name description - description e-mailAddress - e-mail address GN - given name initials - initials L - locality name O - organization OU - organizational unit SN - surname serialNumber- serial number	role mapping rules resource policy rules SSO parameter fields LDAP configuration	certAttr.OU = 'Retail Products Group'

Variable	Description		Examples
	<p>ST - state or province title - title UI - unique identifier Use this variable to check that the user's client has a client-side certificate with the value(s) specified.</p>		
<p>certAttr.altName.<Alt-attr></p>	<p>Subject alternative name value from a client-side certificate where <Alt-attr> may be:</p> <ul style="list-style-type: none"> Email EmailId EmailDomain DNS registeredId ipAddress UPN UPNid UPNDomain fascn fascnAC fascnSC fascnCN fascnCS fascnICI fascnPI fascnOC fascnOI fascnPOA fascnLRC 	<p>role mapping rules resource policy rules SSO parameter fields LDAP configuration</p>	<p>certAttr.altName.email = "joe@company.com" certAttr.altName.ipAddress = 10.10.83.2</p>
<p>certAttr.serialNumber</p>	<p>Client certificate serial number.</p>	<p>role mapping rules</p>	<p>certAttr.SerialNumber = userAttr.certSerial certAttr.SerialNumber = "6f:05:45:ab"</p>


Variable	Description		Examples
	Note that all characters other than [0-9 a-f A-F] are stripped out of a string before comparison with certAttr.SN. Wildcards are not supported.	resource policy rules SSO parameter fields LDAP configurat ion	
certDN	Client certificate subject DN. Wildcards are not permitted.	role mapping rules, resource policy rules	certDN = 'cn=John Harding,ou=eng,c=Company' certDN = userDN (match the certificate subject DN with the LDAP user DN) certDN = userAttr.x509SubjectName certDN = ('cn=John Harding,ou=eng,c=Company' or 'cn=Julia Yount,ou=eng,c=Company')
certDN.<subject-attr>	Any variable from the client certificate subject DN, where subject-attr is the name of the RDN key. Use to test the various subject DN attributes in a standard x.509 certificate.	role mapping rules resource policy rules SSO parameter fields LDAP configurat ion	certDN.OU = 'company' certDN.E = 'joe@company.com' certDN.ST = 'CA'
certDNText	Client certificate user DN stored as a string. Only string comparisons to this value are allowed.	role mapping rules resource policy rules	certDNText = 'cn=John Harding,ou=eng,c=Company'

Variable	Description		Examples
		SSO parameter fields	
certAttr.EKUText	<p>The Enhanced Key Usage field, abbreviated as EKU has 2 components to it. One part of it is the text which is in human readable format and the second part is the OID number which is unique for a given purpose. The user has the flexibility to create rules and realm-based restrictions using either of the two. Format to be given is: EKUText = string or <comma separated string> or string with regular expression. Custom expressions need to be given with the following format: certAttr.EKUText = string or <comma separated string> or string with regular expression.</p>	role mapping rules resource policy rules SSO parameter fields	certAttr.EKUText = "TLS Web Server Authentication","E-mail Protection","TLS Web Client Authentication"
certAttr.EKUOID	Format to be given is:		certAttr.EKUOID=1.3.6.1.5.5.7.3.1,1.3.6.1.5.5.7.3.4,1.3.6.1.5.5.7.3.2

Variable	Description		Examples
	<p>EKUID = to a.b.c.d.e.f.g.h.i or <comma separated list of EKUIDs> or OID with regular expressions This works in both certificate rule as well as custom expressions. Custom expressions need to be given with the following format: certAttr.EKUID = a.b.c.d.e.f.g.h.i or <comma separated list of EKUIDs> or OID with regular expressions</p>		
certIssuerDN	<p>Client certificate- issuer subject DN. This variable works like a standard DN attribute such as CertDN. Wildcards are not permitted.</p>	<p>role mapping rules resource policy rules SSO parameter fields</p>	<p>certIssuerDN = 'cn=John Harding,ou=eng,c=Company' certIssuerDN = userAttr.x509Issuer certIssuerDN = ('ou=eng,c=Company' or 'ou=operations,c=Company')</p>
certIssuerDN.<issuer-attr>	<p>Any variable from the client certificate-issuer subject DN, where issuer-attr is the name of the RDN key.</p>	<p>role mapping rules resource policy rules SSO parameter fields</p>	<p>certIssuerDN.OU = 'company' certIssuerDN.ST = 'CA'</p>

Variable	Description		Examples
certIssuerDNText	Client certificate-issuer subject DN stored as a string. Only string comparisons to this value are allowed.	role mapping rules resource policy rules SSO parameter fields	certIssuerDNText = 'cn=John Harding,ou=eng,c=Company'
defaultNTDomain	Contains the Domain value set in the authentication server configuration when you use AD/NT authentication.	role mapping rules resource policy rules SSO parameter fields	defaultNTDomain=" CORP"
geoLocationCountry	The location from where user should be allowed or denied to login from.	role mapping rules	geoLocationCountry = 'United States' geoLocationCountry = ('United States' or 'Canada')

Variable	Description		Examples
	<p>Note: In case you have a Fresh Installation of ICS/IPS, then it will NOT have UEBA package by default with it. Please add the UEBA package at Behavioral Analysis page before using Adaptive Authentication. In case of Upgrade of ICS/IPS from R7 or earlier to R8 or later, then UEBA package is carried forwarded as is and you can still update it to latest version by uploading new package. You may download latest UEBA package from the Support Site (my.pulsesecure.net).</p>		
<p>group.<group-name></p>	<p>User's group membership as provided by the realm authentication or directory server.</p>	<p>role mapping rules resource policy rules</p>	<p>group.preferredPartner group.goldPartner or group.silverPartner group.employees and time.month = 9 Combination examples: Allow all partners with active status from Monday to Friday but preferred partners Monday through Saturday: (group.partners and time = (Mon to Fri)) or (group.preferredPartners and time = (Mon to Sat))) and userAttr.partnerStatus = 'active'</p>

Variable	Description		Examples
		<p>Only those groups evaluated for role mapping rules are available in the detailed rules (conditions) in the resource policies. We recommend that you use the groups variable instead of group.<group-name>, which is supported only for backwards compatibility.</p>	<hr/> <p> Spaces are not supported, such as, group.sales managers</p> <hr/>
groups	List of groups as provided by the realm authentication or directory server.	role mapping rules	groups=('sales managers')

Variable	Description		Examples
	NOTE: You can enter any characters in the groupname, although wildcard characters are not supported.	resource policy rules SSO parameter fields	
hostCheckerPolicy	Host Checker policies that the client has met.	role mapping rules resource policy rules SSO parameter fields	hostCheckerPolicy = ('Norton' and 'Sygate')
loginHost	Hostname or IP address that the browser uses to contact the Pulse client service.	role mapping rules resource policy rules SSO parameter fields LDAP configuration	loginHost = 10.10.10.10
loginTime	The time of day at which the user submits his credentials. The time is based on system time.	role mapping rules resource policy rules SSO parameter fields	loginTime = (8:00am) loginTime= (Mon to Fri)

Variable	Description		Examples
	<p>NOTE: When using this variable in an SSO parameter field, the variable returns the UNIX string time.</p>		
loginTime.day	<p>The day of month on which the user submits his credentials, where day is 1-31. The time is based on the system time. You cannot use the TO operator with variable.</p>	<p>role mapping rules resource policy rules</p>	loginTime.day = 3
loginTime.dayOfWeek	<p>The day of the week on which the user submits his credentials, where dayOfWeek is in the range [0-6] where 0 = Sunday. The system does not support the TO operator with time.dayOfWeek expressions if you use numbers instead of strings. In other words, "loginTime.dayOfWeek = (2 TO 6)" does not work, but "loginTime.dayOfWeek = (mon to fri)" does work.</p>	<p>role mapping rules resource policy rules</p>	<p>loginTime.dayOfWeek = (0 OR 6) loginTime.dayOfWeek = (mon TO fri) loginTime.dayOfWeek = (1) loginTime.dayOfWeek = 5</p>

Variable	Description		Examples
loginTime.dayOfYear	The numeric day of the year on which the user submits his credentials, where dayOfYear can be set to [0-365]. You cannot use the TO operator with this variable.	role mapping rules resource policy rules	loginTime.dayOfYear = 100
loginTime.month	The month in which the user submits his credentials, where month can be set to [1-12] where 1 = January. You cannot use the TO operator with this variable.	role mapping rules resource policy rules	loginTime.month >= 4 AND loginTime.month <=9
loginTime.year	The year in which the user submits his credentials, where year can be set to [1900-2999]. You cannot use the TO operator with this variable.	role mapping rules resource policy rules	loginTime.year = 2005

Variable	Description		Examples
loginURL	URL of the page that the user accessed to sign in. The system gets this value from the Administrator URLs User URLs column on the Authentication > Signing In > Sign-in Policies page of the admin console.	role mapping rules resource policy rules SSO parameter fields LDAP configuration	loginURL = */admin
networkIf	The network interface on which the user request is received. Possible values: internal, external	role mapping rules resource policy rules SSO parameter fields	sourceIp = 192.168.1.0/24 and networkIf = internal
ntdomain	The NetBIOS NT domain used in NT4 and Active Directory authentication.	role mapping rules SSO parameter fields	ntdomain = jnpr
ntuser	The NT username used in Active Directory authentication	role mapping rules SSO parameter fields	ntuser = jdoe

Variable	Description		Examples
password password[1] password[2]	The password entered by the user for the primary authentication server (password and password[1]) or the secondary authentication server (password[2]).	role mapping rules resource policy rules SSO parameter fields	password = A1defo2z
realm	The name of the authentication realm to which the user is signed in.	role mapping rules resource policy rules SSO parameter fields	Realm = ('GoldPartners' or 'SilverPartners') AND condition will always fail as a user is only allowed to sign in to a single realm in a session.
role	List of all the user roles for the session. In SSO, if you want to send all the roles to back-end applications, use <role sep = ";"> - where sep is the separator string for multiple values. The system supports all separators except " and >.	resource policy rules SSO parameter fields	Role = ('sales' or 'engineering') Role = ('Sales' AND 'Support')

Variable	Description		Examples
sourceIP	The IP address of the machine on which the user authenticates. You can specify the netmask using the bit number or in the netmask format: '255.255.0.0'. Note that you can evaluate the sourceIP expression against a string variable such as an LDAP attribute.	role mapping rules resource policy rules SSO parameter fields	sourceIP = 192.168.10.20 sourceIP = 192.168.1.0/24 and networkIf internal userAttr.dept = ('eng' or 'it') and sourceIP = 10.11.0.0/16 sourceIP = 192.168.10.0/24 (Class C) is the same as: sourceIP = 192.168.10.0/255.255.255.0 sourceIP=userAttr.sourceip
time	The time of day at which the role mapping rule or resource policy rule is evaluated. The time of the day can be in 12-hour or 24-hour format.	role mapping rules resource policy rules	time = (9:00am to 5:00pm) time = (09:00 to 17:00) time = (Mon to Fri) Combination examples: Allow executive managers and their assistants access from Monday to Friday: userAttr.employeeType = ('*manager*' or '*assistant*') and group.executiveStaff and time = (Mon to Fri)
time.day	The day of month on which the user submits his credentials to, where day is 1-31. The time is based on the system time.	role mapping rules resource policy rules	loginTime.day = 3

Variable	Description		Examples
time.dayOfWeek	The day of the week on which the role mapping rule or resource policy rule is evaluated, where dayOfWeek is in the range [0-6] where 0 = Sunday.	role mapping rules resource policy rules	loginTime.dayOfWeek = (0 OR 6) loginTime.dayOfWeek = (1 to 5) loginTime.dayOfWeek = 5
time.dayOfYear	The day of the year on which the role mapping rule or resource policy rule is evaluated. Possible values include: 1-365.	role mapping rules resource policy rules	time.dayOfYear = 100
time.month	The month in which the role mapping rule or resource policy rule is evaluated. Possible values include: 1-12	role mapping rules resource policy rules	time.month >= 9 and time.month <= 12 and time.year = 2004 group.employees and time.month = 9
time.year	The year in which the role mapping rule or resource policy rule is evaluated, where year can be set to [1900-2999].	role mapping rules resource policy rules	time.year = 2005

Variable	Description		Examples
<p>user</p> <p>user@primary_auth_server_name</p> <p>user@secondary_auth_server_name</p>	<p>Pulse client username for the user's primary authentication server (user and user@primary_auth_server_name) or secondary authentication server (user@secondary_auth_server_name).</p> <p>Use when authenticating against an Active Directory server, domain and username.</p> <p>primary_auth_server_name is the name of the primary auth server. If there are spaces or special characters in the name, it can be enclosed in curly brackets. For example, user@{My Primary Auth Server}</p> <p>secondary_auth_server_name is the name of the secondary auth server. If there are spaces or special characters in the name, it can be enclosed in curly brackets. For example, user@{My Secondary Auth Server}</p>	<p>role mapping rules resource policy rules SSO parameter fields</p>	<p>user = 'steve'</p> <p>user = 'domain\\steve'</p>

Variable	Description		Examples
	<p>NOTE: When including a domain as part of a username, you must include two slashes between the domain and user. For example, user='yourcompany.net\joeuser'.</p>		
<p>username username@primary_auth_server_name username@secondary_auth_server_name</p>	<p>Pulse client system username for the user's primary authentication server (username and username@primary_auth_server_name) or secondary authentication server (username@secondary_auth_server_name). If the user is signing in to a certificate authentication server, then the user's Pulse client system username is the same as CertDN.cn. primary_auth_server_name is the name of the primary auth server. If there are spaces or special characters in the name, it can be enclosed in curly brackets. For example user@{My Primary Auth Server}</p>	<p>role mapping rules resource policy rules SSO parameter fields</p>	<p>username = 'steve' and time = mon username = 'steve' username = 'steve*' username = ('steve' or '*jankowski')</p>

Variable	Description		Examples
	<p>secondary_auth_server_name is the name of the secondary auth server. If there are spaces or special characters in the name, it can be enclosed in curly brackets. For example user@{My Secondary Auth Server}</p>		
userAgent	The browser's user agent string.	role mapping rules resource policy rules SSO parameter fields	The browser's user agent string.
userAttr.<auth-attr>	User attributes retrieved from an LDAP, RADIUS, or directory server.	role mapping rules resource policy rules SSO parameter fields	userAttr.building = ('HQ*' or 'MtView[1-3]') userAttr.dept = ('sales' and 'eng') userAttr.dept = ('eng' or 'it' or 'custsupport') userAttr.division = 'sales' userAttr.employeeType != 'contractor' userAttr.salaryGrade > 10 userAttr.salesConfirmed >= userAttr.salesQuota Negative examples: userAttr.company != "Acme Inc" or not group.contractors not (user = 'guest' or group.demo) Combination examples: Allow executive managers and their assistants access from Monday to Friday:

Variable	Description		Examples
			<p>userAttr.employeeType = ('*manager*' or '*assistant*') and group.executiveStaff and time = (Mon to Fri)</p> <p>Allow all partners with active status from Monday to Friday but preferred partners Monday through Saturday:</p> <p>((group.partners and time = (Mon to Fri)) or (group.preferredPartners and time = (Mon to Sat))) and userAttr.partnerStatus = 'active'</p>
userDN	<p>The user DN from an LDAP server (not applicable to Active Directory auth server with ldap group lookup). If the user is authenticated by the LDAP server, then this DN is from the authentication server; otherwise, the DN comes from the realm's Directory/Attribute server.</p>	<p>role mapping rules resource policy rules</p>	<p>userDN = 'cn=John Harding,ou=eng,c=Company' userDN = certDN</p>
userDN.<user-attr>	<p>Any variable from the user DN, where user-attr is the name of the RDN key.</p>	<p>role mapping rules resource policy rules</p>	<p>Any variable from the user DN, where user-attr is the name of the RDN key.</p>

Variable	Description		Examples
		SSO parameter fields	
userDNText	User DN stored as a string. Only string comparisons to this value are allowed.	role mapping rules resource policy rules SSO parameter fields	userDNText = 'cn=John Harding,ou=eng,c=Company'

Custom Variables and Macros

Custom variables, like system variables, are name-value pair tags that you can use when defining role mapping rules, resource policy rules and SSO parameter fields.

Custom variables are created in the Server Catalog (for example, **Authentication > Auth Server > Name > Settings**) by using a predefined macro on a system variable. Available macros are:

- REGMATCH - Matches a regular expression pattern against a string text.
- APPEND - Appends a text string to another text string.
- DAYSDIFF - Calculates the difference between two dates.



These macros are located under Variable Operators in the Variables tab of the Server Catalog window.

A custom variable name is a dot-separated string. Each component can contain characters from the set [a-z A-Z 0-9 _] but cannot start with a digit [0-9]. Custom variable names are case-insensitive.

Custom variables are referenced as **customVar.<variableName>**. For example, if you create a custom variable with the name **check-prefix**, you reference this custom variable as **customVar.check-prefix**.

append

Field	Description
Syntax	APPEND (attr, TextString) APPEND (attr, attr2)
DescriptionS	Append a text string to an attribute or append an attribute to another attribute and store the resulting string in the custom variable.
Options	attr -System variable of type string. TextString -Quoted ASCII string. attr2 -System variable of type string.
Output Fields	Returns a String value. If no match is found, returns an empty string. If the system variable is multivalued, the custom variable is also multivalued and uses the same order as the system variable.
Sample Output	APPEND (userName, "@pulsesecure.net") In this example, the string "@pulsesecure.net" is appended to the userName value.

daysdiff

Field	Description
Syntax	DAYSDIFF (attr, timeformat)
Description	Calculates the number of days between the attribute and the current time.
Options	attr -System variable of type string. timeformat -Output time format. Valid values are: UTC, TIMET, MMDDYYYY
Output Fields	Returns an Integer value.
Sample Output	DAYSDIFF (certAttr.validUpto, UTC) In this example, calculate the difference in days between the current time and the value of certAttr.validUpto and express the time in UTC (Coordinated Universal Time).

regmatch

Field	Description
Syntax	REGMATCH (attr, regex, groupingNumber)
Description	Match the regular expression pattern against an attribute and store the result in the custom variable.
Options	<p>attr-System variable of type string.</p> <p>regex-Quoted string containing the regular expression to be applied to the attr option.</p> <p>groupingNumber-The group value to assign to the custom variable.</p>
Additional Information	The regular expression supports the Perl Compatible Regular Expressions (PCRE) syntax. A grouping (capture buffer) in the regex pattern can also be used to define a custom variable.
Output Fields	Returns a String value. If no match is found, returns an empty string. If the system variable is multivalued, the custom variable is also multivalued and uses the same order as the system variable.
Sample Output	REGMATCH (mailId, "^(.*)@pulsesecure.net\$", 1) In this example, a mailId of myName@pulsesecure.net creates a custom variable with value "myName".

Specifying Fetch Attributes in a Realm

To facilitate the support for various parameterized settings in user roles and resource policies, you have the ability to specify additional fetch attributes. The system stores the fetch attributes when users log in so that you can use them in parameterized role or resource policy definitions.

The system pulls all the attributes that are currently stored in the Sever Catalog for the user's authentication or authorization LDAP server. So, make sure to add the LDAP user attributes that are used in role or resource policy definitions in the LDAP Server Catalog first.

When a user logs in, the system retrieves user attributes that are referenced in the role mapping rules plus all of the additional attributes referenced in the Server Catalog and stores all these values. Note that this should not incur a significant performance overhead because all the user attributes are retrieved in one single LDAP query.



When you substitute variables, such as in IP/Netmasks or hostnames, the values in the session are appropriately converted into the data type that is required by the particular application definition.

Specifying the homeDirectory Attribute for LDAP

You can create a bookmark that automatically maps to a user's LDAP home directory. You can accomplish this using the LDAP attribute homeDirectory. You need to configure a realm that specifies the LDAP server instance as its auth server, and you need to configure role-mapping rules and a bookmark that points to the LDAP homeDirectory attribute.