ivanti

Neurons for Zero Trust Access Tenant Admin Guide

Copyright Notice

Contents

END USER LICENSE AGREEMENT

The Ivanti product that is the subject of this technical documentation consists of (or is intended for use with) Ivanti software. Use of such software is subject to the terms and conditions of the End User License Agreement (“EULA”) posted at https://www.ivanti.com/company/legal/eula. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

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
italic text	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.
italic text	Identifies a variable.

Convention	Description
[ ]	Syntax components displayed within square brackets are optional. 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.

Code Block

Following is an example of Python based code block in the html documentation:

```
def some_function():
interesting = False
print 'This line is highlighted.'
print 'This one is not...'
print '...but this one is.'
```

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.

About This Guide

To see an overview of the software and workflows needed to get started quickly, read the *Neurons for Zero Trust Access: Getting Started Guide*.

This guide takes you through the configuration and use of your Ivanti Neurons for Zero Trust Access (nZTA) service.

It covers the following subject areas:

- An overview of nZTA network architecture and concepts. See Overview of Ivanti Neurons for Zero Trust Access.
- Logging into the nZTA Controller as a Tenant Admin. See Logging in as a Tenant Administrator.
- Setting up one or more Gateways. See Working with Gateways.
- Defining user authentication methods and policies, see Working with User Authentication.
- Creating device policies. See Creating Device Policies and Device Policy Rules.
- Adding applications to your Gateways. See Working with Applications and Application Groups.
- Configuring secure access policies. See Creating/Editing Secure Access Policies.
- Enrolling desktop and mobile devices with the nZTA Controller. See Enrolling Ivanti Secure Access Client.
- Updating the client software and ESAP definition on Ivanti Secure Access Client, see Upgrading Ivanti Secure Access Client.
- Using analytics to monitor your nZTA platform and services. See Using the Insights Menu to Monitor User Activity and Service Usage.

What's New

Version 22.7R1

Configurable MTU size for gateways

Tenant admin can now define MTU size for ZTA gateways depending on their requirements and underlying network infrastructure.

For details, see:

- "Adding a VMware vSphere Gateway" on page 244
- "Adding an Amazon Web Services Gateway" on page 255
- "Adding an Azure Gateway" on page 268
- "Adding a KVM Gateway" on page 288
- "Adding a GCP Gateway" on page 307
- "Adding an Oracle Gateway" on page 330

Password Strengthening for Local Authentication Server

The local authentication server has stronger password restrictions.

For details, see "Workflow: Creating a Local Authentication Policy" on page 102.

Renewed ZTA IDP metadata in release 22.7R1

To ensure continued compatibility, download the renewed ZTA IDP metadata from the ZTA tenant application configuration page and subsequently apply the updated information to the SaaS SAML SSO configurations.

Version 22.6R1.2

iOS supporting Ivanti NMDM integration with nZTA

Ivanti Neurons for MDM provides compliance check and simplified onboarding experience for nZTA end users connecting via mobile. For details, see "Integrating Ivanti Neurons for MDM with nZTA" on page 71.

Hardened custom sign-in policies and login URLs

As part of hardening custom sign-in policies and login URLs, the following changes are implemented:

- Instead of requiring administrators to configure enrollment policies, administrators will only need to configure user policies. As a default, all configured user policies support enrollment.
- Single SAML authentication server for user authentication and enrollment.

For details, see:

- "Workflow: Creating a SAML Authentication Policy With Azure AD" on page 113
- "Workflow: Creating an Authentication Policy for On-Premises ICS SAML" on page 145
- "Workflow: Creating a SAML Authentication Policy for Okta" on page 164
- "Workflow: Creating a SAML Authentication Policy for Ping Identity" on page 182
- "Workflow: Creating a Local Authentication Policy" on page 102

Version 22.6R1

Oracle Cloud Platform support for nZTA Gateway

nZTA Gateway now supports deployment on Oralce Cloud Platform (OCI). For details see "Workflow: Creating a Gateway in Oracle Cloud Platform" on page 328.

Launching the Windows Edge/Webview2 browser

In a typical enrollment, upon successful authentication to the Controller, Ivanti Secure Access Client automatically shows the end-user portal applications page through a Windows Edge/Webview2 browser. This feature is supported with ISAC client version 22.6R1. For details. see "Enrolling a Windows Device" on page 510.

Reusable custom icon to associate with application

The create application page provides an option to upload your own icon, which can be re-used to associate with more than one application. For details. see "Working with Applications and Application Groups" on page 477.

Admin experience enhancements to L4, Gateway Logs, and Logs Tables in terms of selection and resizing, pagination, text copy/paste, confirmation actions across all pages.

The following list shows the enhancements to L4, Gateway Logs, and Logs Tables.

- Column resizing across nZTA pages
- Cell content copy text from Table
- Pagination across nZTA pages
- Minimum number of columns in all the tables in L4 dashboards
- Enhancement to Advanced Filter

For details. see "Viewing Detailed Logs for a Chart" on page 616, and "Filtering the Logs" on page 689.

Simplifying Devices section and better correlation of data for Device Rules and Policies and Device Insights with a new workflow

Admin experience is enhanced by simplifying the device rules and policies, and global device preferences. For details, see "Creating Device Policies" on page 443, "Setting Global Device Preferences" on page 469, and "Viewing Currently Enrolled User Devices" on page 666.

Version 22.5R1.2

Suppress EUP Auto Launch

Allows Admin to suppress the auto launch of the End User portal. This option is enabled by default and works with ISAC 22.5R1 and later. For details, see "Enrolling Ivanti Secure Access Client" on page 508.

Version 22.5R1

Admin Access Control Based on Location, Host Checker, and Network

Checks the Admin's device geographic location/network/host checker compliance for admin sign-in policy before providing access to admin login. For details, see "Configuring Default Device Policy for Users" on page 448.

Enhancements to Anomalies and Non-compliance L4 Drill Down logs

The Anomalies L4 table now includes MAC Address and Source IP Address columns.

The Non-compliances L4 table now includes Acknowledged, Non-compliant Policy Type, Non-compliance Policy reason, MAC Address and Source IP Address columns.

For details, see "Using the Active Anomaly, Connected Clients Version, and Non-Compliance Charts" on page 613.

Log export options to the admin from Gateway and L4 (drill-down view) logs

In any of the L4 pages, export the displayed log as a CSV or JSON text file, or create schedules to set up log export jobs. For details, see "Viewing Detailed Logs for a Chart" on page 616.

Exporting logs from L4 (drill-down view) logs and Gateway logs

For details, see "Exporting Logs" on page 691.

Gateway Creation Config UI Simplification

Create nZTA Gateway and Create nZTA Gateway Group options are grouped under Create. For details, see "Adding a VMware vSphere Gateway" on page 244.

Acknowledge non-compliance in the non-compliance info panel

Acknowledge individual non-compliances and remove them from the active total. Filter on acknowledged, unacknowledged (active), or all non-compliances. For details, see "Using the Summary Ribbon" on page 597.

Create another application option in Applications page

On Create Application page, admin can choose to continue to create one more application. For details, see "Working with Applications and Application Groups" on page 477.

Version 22.4R3

Role Based Access Control for Admin Users

With Role-based access control (RBAC), organizations can easily add admins and assign them specific roles, with differing levels of access to the nSA Admin Portal. In addition to an existing set of default roles, Administrators can now create custom granular roles for specific functions within the nSA admin portal.

For details, see "Role-based Access Control for Admin Users" on page 99

HTTP Proxy Support

Support Proxy configuration in gateway to connect to nZTA.

For details, see:

"Adding a VMware vSphere Gateway" on page 244

"Adding an Amazon Web Services Gateway" on page 255

"Adding an Azure Gateway" on page 268

"Adding a KVM Gateway" on page 288

"Adding a GCP Gateway" on page 307

Version 22.4R1

Applications and Application Groups UI change

Group together multiple applications for which a single secure access policy is required, For details, see "Working with Applications and Application Groups" on page 477 and "Adding Application Groups to the Controller" on page 488.

nZTA Gateway Connection Control for Trusted Networks

nZTA Gateway can sometimes be bypassed so that users can connect directly to specific applications. For example, you might want users to bypass nZTA for a specific application if they are connected directly to your trusted corporate network. nZTA Gateway tunnel creation will be bypassed on the endpoint since resource access will go through the physical interface.

For details, see "Configuring nZTA Gateway Connection Control for Trusted Networks" on page 438

Gateway Re-registration

nZTA Gateway can now be re-registered in case if the Gateway Registration was not successful and can edit gateway configuration parameters. On registration failures, admin can trigger the registration manually along with the current debugging options such as networking tools, reboot etc. You can also regenerate and download the gateway init config from the controller admin interface as when required. The Admin can also use Registration error report, which provides insight about the registration failure and suggest solutions to overcome it.

Limitation: Azure and KVM does not allow the user to update configuration after the gateway is deployed. So, if any config update is needed in Azure or KVM Gateways (nZTA) ,we need to redeploy the nZTA Gateway.

For details, see "Re-registering a VMware vSphere Gateway" on page 251, "Re-registering an Amazon Web Services Gateway" on page 264 and "Re-registering a GCP Gateway" on page 310.

Location/Network rule support in default device policy

Location/Network policy based enforcement can be applied for any user policy. For details about enforcing the policy on default device, see "Configuring Default Device Policy for Users" on page 448.

For details about creating these policies, see "Options for Location Rules" on page 460 and "Options for Network Rules" on page 462.

Overview of *Ivanti Neurons for Zero Trust Access*

Securing a Diverse Application Infrastructure

To enhance an enterprise's ability to protect its systems, services, and data, IT teams can adopt a *zero-trust* network security architecture. The zero-trust model works on the principle that no single client entity is automatically trusted, regardless of whether they are inside or outside the organizational network perimeter. Each client connection must have its identity verified by a trusted service or policy before access is granted.

As organizations move away from traditional application deployment models and begin to adopt hybrid and multi-cloud service-based infrastructures, the notion of zero-trust must be a extended to encompass the means to centrally manage application access and security for all users across the enterprise, regardless of device type and location.

For example, an enterprise with SaaS, 3rd-party cloud, and on-premise applications might need to define end-to-end secure access policies through a combination of:

- User authentication method (Active Directory, LDAP, local authentication)
- Device compliance policy (operating system, antivirus definition, root access)
- Infrastructure type (AWS/Azure, on-premise vSphere, KVM)
- Application type (SaaS applications such as O365/SFDC, Cloud apps, on-premise apps)
- Additional security postures (such as location, time, usage patterns)

In addition, this model gives rise to complex requirements for visibility and compliance reporting. At any time, an organization might need to see the following types of information:

- The current number of compliant users or devices
- Application usage levels both in the cloud and on-premise
- Visibility of any abnormal or compliance issues

Hyper-Converged Zero-Trust Access

Ivanti Neurons for Zero Trust Access (*nZTA*) addresses the requirement for an enterprise-level zero-trust secure access management platform. It offers the following features and benefits:

- True zero-trust access as a service for hybrid and multi-cloud
    - Authentication and authorization before any access
- Separation of control and data planes
    - All user data goes through only the relevant *nZTA Gateway* deployed in the organization's datacenter or virtual private cloud
- Dark cloud to reduce the attack surface
    - *nZTA Gateway*s are not accessible by the client end-point until authorized to do so by the *Controller*
    - Client end-points never access the applications directly; only through an encrypted data tunnel via the *nZTA Gateway*
- Smooth migration
    - Users continue to use existing VPN connection profiles through *Ivanti Secure Access Client*, while incorporating new *nZTA* profiles to access your managed application infrastructure

An Overview of *Ivanti Neurons for Zero Trust Access*

Ivanti Neurons for Zero Trust Access (*nZTA*) is a cloud-based SaaS (software as a service) application that forms part of the *Ivanti Neurons for Secure Access* family. It provides fully-managed zero-trust authentication and access control for an organization's application infrastructure.

nZTA is developed according to the principle of SDP (Software Defined Perimeter). Where traditional network-based security architectures use firewalls at the network perimeter to limit access to resources, SDP is based on a need-to-know policy-driven model. Client devices are verified and authorized before access to applications is granted. Application infrastructure cannot be detected remotely, and has no visible DNS information or exposed IP addresses. This protects networked resources from many common network-based attacks.

nZTA enables system administrators to define end-to-end authorization and authentication policies that control application visibility, access, and security for all users and their devices. Administrators can deploy applications of any type in a variety of hybrid cloud and on-premise datacenter environments, with *nZTA* providing users seamless secure access to only those applications for which they are authorized to use, regardless of geographic location and client device platform.

nZTA facilitates authorization policy enactment for any combination of users, devices, applications, and infrastructure.

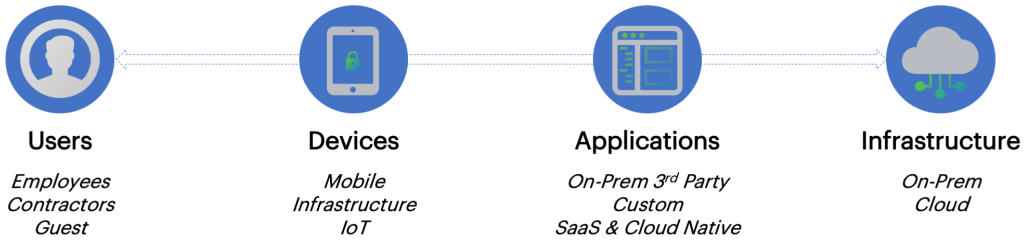

Four-dimensional authentication policy control

Deploying and Using *nZTA*

A *nZTA* deployment consists of a *Controller* service and one or more application *Gateways* positioned at each location an organization hosts its resources and applications. This might be in a public or private cloud, within a datacenter, or inside a virtual host environment.

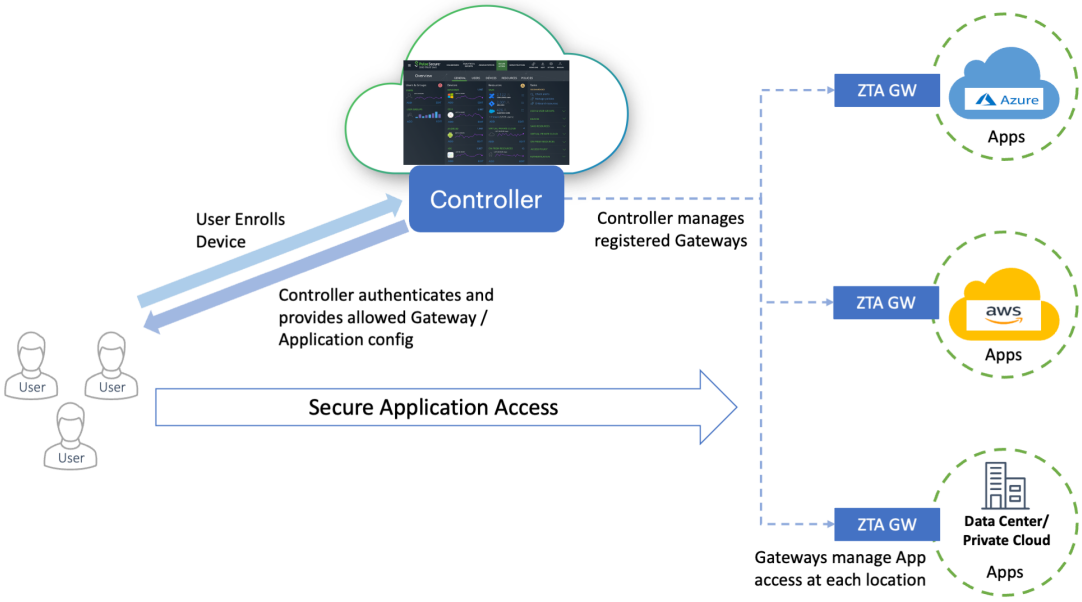

The topology of a *nZTA* deployment

End users use *Ivanti Secure Access Client* to authenticate with the *Controller*, which runs an Authentication, Authorization and Accounting (AAA) Service. The *Controller* then enables direct and encrypted communication between the user and the *nZTA Gateway*s that protect the user's authorized resources. This mechanism avoids the general exposure of public IP addresses, and separates the control plane from the data plane.

To learn more about using *Ivanti Secure Access Client* with *nZTA*, see Enrolling Mobile/Desktop Clients and Using Ivanti Secure Access Client with nZTA.

After you have defined your secure access policy, or policies, the *Controller* automatically synchronizes the configuration out to your *nZTA Gateway*s.

System administrators use the *Controller* to configure, provision, and monitor an enterprise's application infrastructure. To achieve this, *Ivanti* provides a *Tenant Admin Portal* on the *Controller*. This portal facilitates the following actions:

- Configure a custom domain at which your *nZTA* service is reached
- Define and deploy *nZTA Gateway*s
- Register authentication services and policies
- Publish resources and applications
- Create authorization rules and policies for users and devices
- Enroll users
- Perform management and real-time analysis of traffic flow and application access across the enterprise
- Manage your Tenant Admin Portal settings

To login to the Tenant Admin Portal, see Logging in as a Tenant Administrator.

Using a Custom Domain

Your *nZTA* tenant subscription is deployed, by default, as a unique endpoint at a domain provided by *Ivanti*. Enrollment and sign-in endpoints are then configured at this FQDN through *user authentication policies*.

Should you require it, you can provision a custom domain to be used in place of the *Ivanti*-provided domain.

IMPORTANT: It is essential to determine whether a custom domain is required at the outset of your subscription, and particularly before you have deployed any Gateways or enrolled any users.

To learn more, see Specifying a Custom Domain.

Deploying Gateways

nZTA requires you to set up Gateways to manage access to your applications and resources. You deploy one or more Gateway instances at each location your applications are hosted, whether at a physical datacenter, a private or public cloud-based service, or some hybrid combination. Each Gateway communicates with the *Controller* to ensure that access requests are authorized.

nZTA supports two main Gateway types, depending on your subscription:

- *nZTA Gateway*s
- Ivanti Connect Secure (ICS) Gateways

This guide describes how to install and configure *nZTA Gateway*s. For details pertaining to an ICS Gateway, refer instead to the "ICS Gateway Onboarding Guide" in the *nSA* documentation portal.

A *nZTA Gateway* is a virtual machine that operates as a headless instance at the perimeter of the logical network used by your datacenter or cloud. A *nZTA Gateway* must be contactable by the *Controller*, other *nZTA Gateway* instances at the same location, and the applications that reside there. *nZTA* supports *nZTA Gateway*s running on Microsoft Azure, Amazon Web Services (AWS), and VMware vSphere. *Ivanti* provides template images for each environment.

To obtain the latest *nZTA Gateway* virtual machine template images, see the Release Notes.

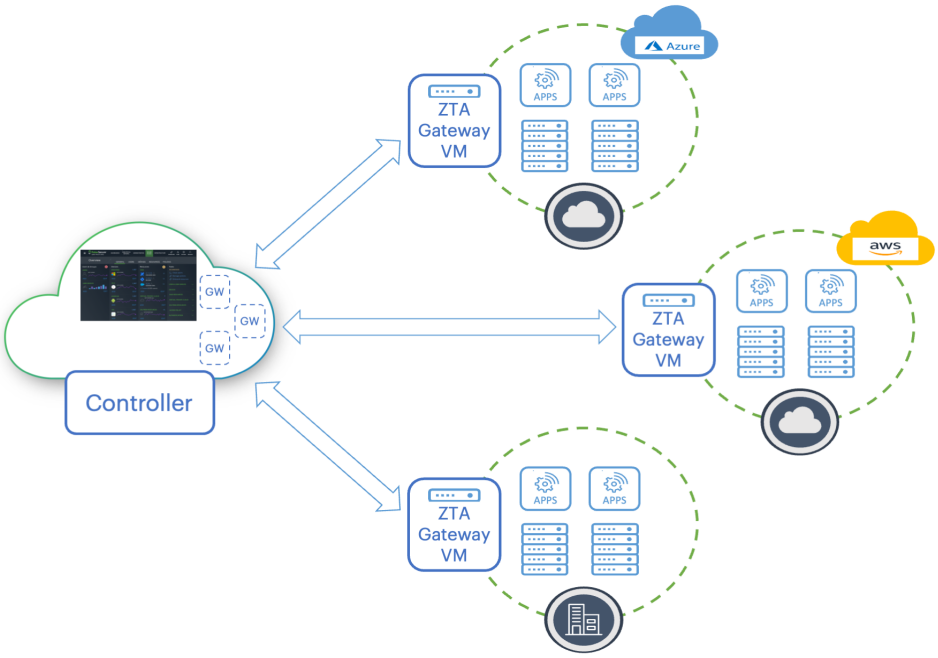

Deploying *nZTA Gateway*s to your cloud and datacenter locations

The process of registering a *nZTA Gateway* identity with the *Controller* produces a package of settings known as a *Gateway definition* that you apply to the Gateway virtual machine during deployment. These settings enable the *nZTA Gateway* to establish secure communication with the *Controller*.

Make sure the *nZTA Gateway* instance does not exist prior to registration with the *Controller*. The Gateway definition file is designed to be applied to a new virtual machine instance at deployment time only.

To register a new *nZTA Gateway* with the *Controller*, use the Tenant Admin Portal. The *Controller* provides a step-by-step workflow to define your *nZTA Gateway*s, requiring basic identification and networking details. After you complete this process, you can download the definition file ready for insertion into your *nZTA Gateway* virtual machine at deploy time.

A Gateway definition file is valid for 24 hours. If this period expires, you must replace the *nZTA Gateway* to generate a new definition file.

Then, when you create the *nZTA Gateway* virtual machine instance in your cloud or on-premise management console, you apply the definition file when requested by the template. A newly launched *nZTA Gateway* instance registers itself with the *Controller*, and any subsequent policy changes made on the *Controller* are automatically synchronized out to all *nZTA Gateway*s.

To learn more about the process of creating and managing your *nZTA Gateway*s, including network interface and subnet requirements, see Working with Gateways.

Using Gateway Groups for High Availability

nZTA provides the ability to enhance your application delivery by allowing the deployment of a group of two or more *nZTA Gateway* instances in front of the same applications. Depending on the approach taken, this can provide improvements for:

- **Availability**: If one Gateway becomes overloaded or unresponsive, traffic is routed to another Gateway in the group.
- **Throughput**: Balancing the load between all Gateways in the group to take full advantage of their capacity.
- **Load partitioning**: A non-balancing approach to distributing traffic between Gateways in the group.
- **Latency reduction**: Routing an end user's traffic to the Gateway nearest to them.

A *nZTA* deployment provides two methods to control how traffic is routed:

1. The *Controller* chooses which Gateway to use, see Using Controller for Gateway Selection.
2. You deploy a load-balancer to route traffic to the Gateways, see Using an Administrator-Deployed Load-Balancer.

The *Controller* uses the concept of *Gateway Groups* to manage each of these deployment models. A Gateway Group contains Gateways that are considered part of a single high-availability group, see Using Gateway Groups in your Secure Access Policies.

To use active/standby and geo-proximity traffic distribution features with your *nZTA* deployment, your end-users must be running *Ivanti Secure Access Client* versions supported for use with *nZTA* version 21.3.2 or later. To learn more about supported client versions, see the *Release Notes*.

Using *Controller* for Gateway Selection

To allow the *Controller* to select a Gateway, make sure your Gateway Group has no load-balancer defined as part of its configuration.

In this scenario, no additional infrastructure is required.

When Gateways in the group are deployed across different locations, the *Controller* selects the Gateway that is geographically closest to the end-user, based on the end-user's IP address. When Gateways in the group are configured with the same location, the end-user is directed to the Gateway that is alphabetically first by name. If a gateway fails, the end-user is automatically directed to use the next available Gateway in the group by the same criteria.

This mode provides improved availability over a standalone Gateway. In addition, when Gateways are deployed across multiple locations, it can partition load and reduce latency for end-users.

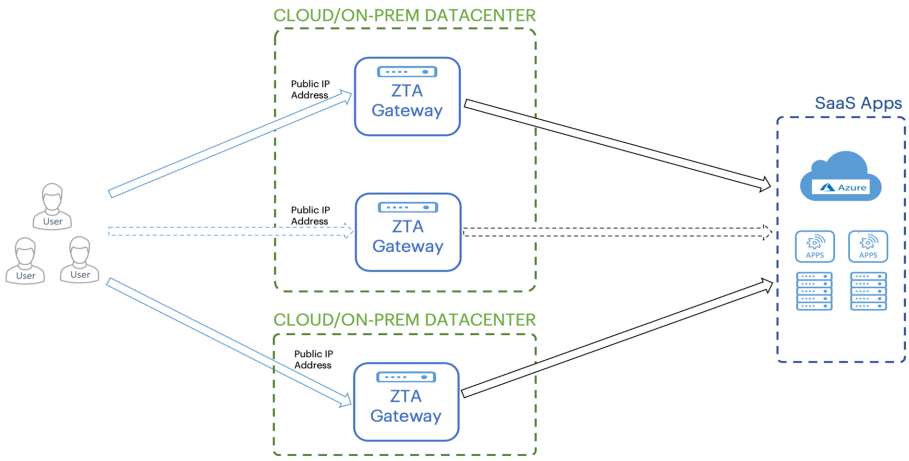

Using *Controller* Gateway Selection with Gateways over multiple locations

Using an Administrator-Deployed Load-Balancer

An administrator can deploy a load-balancer to distribute traffic between the Gateways in the group through the following scenarios:

- As an *in-path* load-balancer to provide more rapid failure detection, and to load-balance traffic across Gateways when they are deployed in the same location to maximize throughput. An in-path load-balancer does not provide for latency reduction as all the Gateways in the group are deployed in the same location.

  For on-premise Gateway deployments, *Ivanti* recommends **Pulse Secure Virtual Traffic Manager (vTM)**. For cloud-based Gateway deployments, you can again use vTM or the built-in load balancing capabilities of AWS or Azure. To learn more about vTM, contact *Ivanti* Technical Support or visit www.ivanti.com. For more details on available load balancers in AWS or Azure, see the relevant product documentation.

  When using vTM for load-balancing, *Ivanti* recommends you configure the following settings in the vTM Admin UI to ensure proper fail-over functionality:

  1. Click **Services > Virtual servers** and edit your designated HTTPS virtual server. In the virtual server configuration, click **Protocol Settings > TCP Connection Settings** and enable **close_with_rst**. This setting closes connections from clients with a RST packet rather than a FIN packet and avoids the TIME_WAIT state.
  2. Click **Services > Pools** and edit your designated HTTPS pool. In the pool configuration, enable the following settings:
     - Click **Protocol Settings > TCP Pool Settings** and enable **node_connclose**. This setting closes all connections to a node if vTM detects that the node has failed.
     - Click **Protocol Settings > TCP Protocol Options** and enable **node_close_with_rst**. This setting closes connections from clients with a RST packet rather than a FIN packet and avoids the TIME_WAIT state.

  You deploy an in-path load-balancer in front of your Gateway instances and raise a public IP address on its front-end interface. You then configure the load balancer with a pool of nodes representing the external IP addresses of your Gateway instances. End users connect to the load-balancer, which in turn forwards on the connection to one of the available Gateways based on the load balancing algorithm you choose (such as round-robin). The load balancer can be configured with a health monitor to perform connectivity checks with your Gateways.

To use an in-path load-balancer, enter the IP address or FQDN of the load-balancer's public interface in the **load-balancer** field of the Gateway Group configuration (see Working with Gateways).

nZTA supports Gateway health monitoring using TCP port 443.

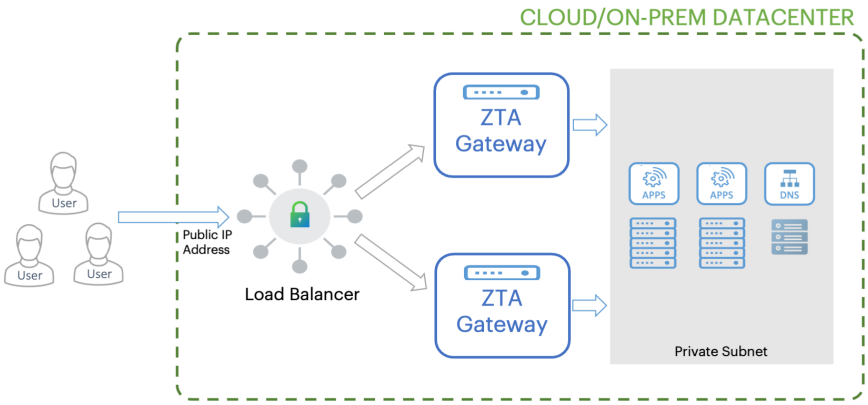

Deploying a pair of Gateways for high availability

This diagram demonstrates high availability based on a pair of Gateways. You can add further Gateway instances according to the expected load on the applications at that location.

- As a *global* load-balancer to distribute traffic between Gateway instances based on client proximity to the Gateway.

  The global load-balancer should be configured to manage an FQDN and resolve it to the public IP address of the optimal Gateway instance. To use the global load-balancer, add the FQDN managed by the global load-balancer to the **load-balancer** field of your Gateway Group configuration.

Using Gateway Groups in your Secure Access Policies

To implement high availability, you configure a *Secure Access Policy* to use a Gateway Group rather than an individual Gateway. For all policy updates, the *Controller* automatically synchronizes changes with all Gateways in the group.

Note the following when using Gateway Groups:

- A Gateway is available to be assigned to a Secure Access Policy as part of a Group, or as an individual instance, but not both. That is, a Gateway that is added to a Gateway Group is no longer available to be assigned to a Secure Access Policy by itself.
- A Gateway must be registered and connected before it is available to be added to a Gateway Group.
- Where the *Controller* is used for Gateway selection, all Gateways in an affected group must have a public IP address configured.
- Gateway instances use different certificates for mTLS communication when part of a Gateway Group. When you add a Gateway to a Gateway Group, this triggers a task to renew the Gateway certificate to match that used by the Group. Conversely, if you remove a Gateway from a Gateway Group, this again triggers a task to renew the Gateway certificate to ensure validity as an individual instance.
- A removed Gateway does not retain any user or device policies previous synchronized to it by the *Controller*.

To learn more about configuring a Gateway Group, see Working with Gateways.

Using Network Load Balancer for High Availability in AWS

If you plan to use Network Load Balancer (NLB) to implement high availability for a group of Gateways deployed in AWS, *Ivanti* recommends you configure a **Target Group** with a *Target Type* of "IP":

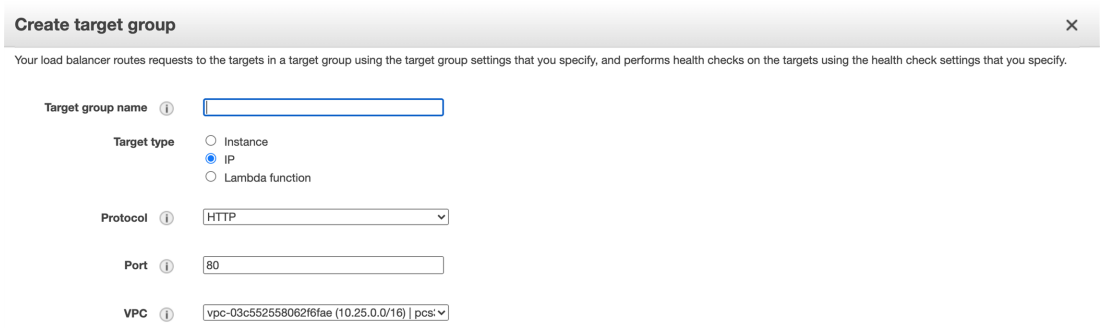

Setting Target Type in a Target Group

Defining User Authentication

nZTA provides user authentication through **authentication policies**. Policies cover the primary scenarios of user enrollment, user sign-in, and administrator sign-in, and the *Controller* includes built-in policy definitions for all three. *nZTA* additionally allows you to define your own custom enrollment and sign-in policies to facilitate specific authentication paths for different groups of users or parts of your organization.

An authentication policy defines the application of an **authentication method** for a specified access URL. You create methods based on your user authentication requirements and assign them to the appropriate policies. As a user connects to the access URL defined in a policy, the associated method is employed to authenticate the user device.

nZTA facilitates *Multi-Factor Authentication* (MFA) through the configuration of an optional secondary authentication method in a policy. When MFA is deployed, a connecting user would need to satisfy the requirements of both primary and secondary authentication methods before access is granted. *nZTA* allows the use of *Local authentication* and *Time-based One Time Password (TOTP)* as secondary methods.

nZTA also provides for the definition of **user rules** and **user groups**. User rules act as filters and define the basic criteria by which users' credentials must match in order for authentication to proceed. User groups encapsulate an authentication policy with one or more user rules to provide a complete user authentication definition for your *Secure Access Policy*.

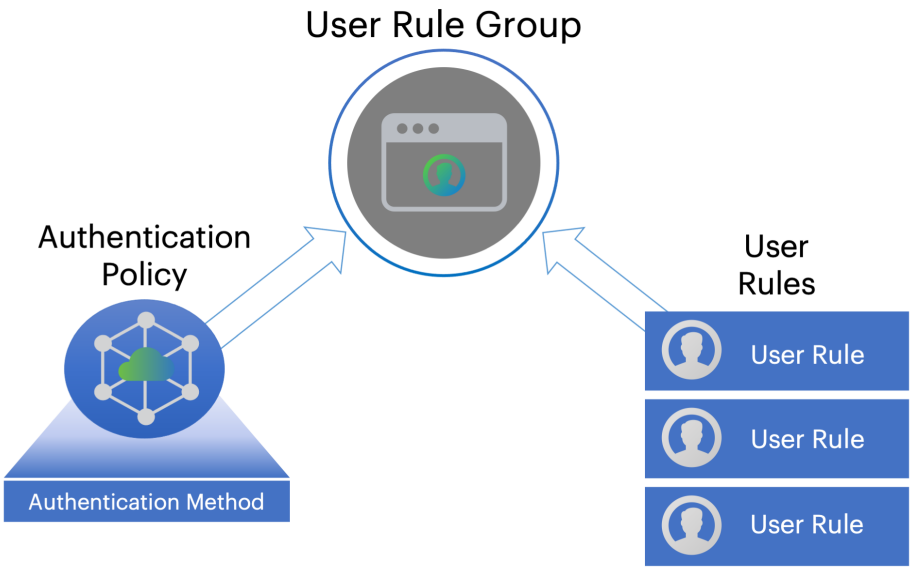

The relationship between user groups, rules, authentication policies and methods

Ivanti recommends you define your authentication methods and user rules first. Then, create or update authentication policies for user enrollment and sign-in based on the authentication methods that you want to apply in that scenario. Finally, create a user group based on your policy and add any applicable user rules.

- To read more about authentication methods, see Creating User Authentication Methods.
- To read more about authentication policies, see Using User Authentication Policies.
- To read more about user rules and groups, see Defining User Authentication Rules and Groups.

Creating User Authentication Methods

nZTA supports user authentication through the following methods:

- local authentication
- SAML authentication (through Azure AD or a custom defined service)
- Time-based One Time Password (TOTP) authentication

Local Authentication

nZTA can operate a local authentication policy based on a defined list of users held internally in the *Controller*. User authentication requests are compared against this list to determine access rights. You create and manage this list manually through the Tenant Admin Portal.

nZTA supports the use of Local Authentication as a *primary* or *secondary* authentication mechanism.

To learn more about defining local authenticators in *nZTA*, see Workflow: Creating a Local Authentication Policy.

SAML Authentication

`SAML (Security Assertion Markup Language)` is an open-standard XML-based message framework, used to exchange information concerning the authentication and authorization of user agents attempting to access an organization's web services. SAML is often used for web browser single sign-on (SSO) services.

SAML deployments typically involve two participant types:

- Service Provider (SP): provides the resources to be protected
- Identity Provider (IdP): performs the authentication and authorization checks required by your resources

nZTA can function as a SAML SP to control user access to your application infrastructure, using an externally-defined service (for example, Azure AD) as the IdP. Access is permitted only when a valid SAML token is presented by the requesting user agent, who has in turn obtained the token from the IdP.

nZTA supports the use of SAML in Multi-Factor Authentication as a *primary* authentication mechanism only.

To configure *nZTA* to use a SAML authenticator, you must establish a trusted SAML connection between *nZTA* and a suitable identity provider. To create the connection, *nZTA* and the identity provider must both be configured with authentication data identifying each other.

In other words, when you configure an identity provider application to operate as a SAML IdP, you obtain an *IdP metadata* package file containing configuration details suitable for uploading to *nZTA* (as the SP). After you upload the IdP metadata to *nZTA*, *nZTA* in return generates its own *SP configuration metadata* package file for uploading back to the IdP.

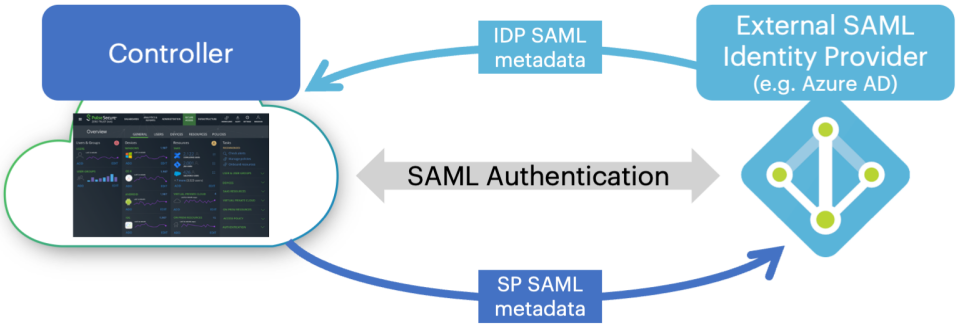

Configuring *nZTA* with a SAML authentication method

This process of authentication metadata-sharing establishes the IdP as a trusted authenticator within *nZTA*. The IdP, in turn, acknowledges *nZTA* as a trusted SP.

This guide provides information on configuring trusted connections to supported SAML authenticators. For more detailed instructions, or for deployment scenarios not covered by this guide, refer to the documentation provided by your IdP vendor.

nZTA uses independent authentication policies to handle *user enrollment* and *user sign-in* processes. When configuring SAML authentication, you typically create two separate SAML applications in the IdP to handle each process individually. This means that for each authentication policy you must configure a separate *nZTA* authentication method and, for each method, perform the IdP/SP metadata handshake to establish a trusted connection with the associated SAML application. *Ivanti* recommends you configure your IdP SAML applications and obtain the corresponding metadata files before you configure your *nZTA* authentication methods.

To learn more about configuring SAML authenticators in *nZTA*, see Working with User Authentication.

TOTP Authentication

Time-based One Time Password (TOTP) is defined in RFC6238 as an authentication mechanism where a one-time password (also known as a token) is generated by an authentication server and client from a shared secret key and the current time. This is implemented in *nZTA* with the *Controller* acting as the TOTP authentication server.

Any third-party TOTP applications (for example, Google Authenticator, Microsoft Authenticator, and so on) available on the mobile and desktop client platforms generate TOTP tokens.

nZTA supports the use of TOTP in Multi-Factor Authentication as a *secondary* authentication mechanism only.

To learn more about defining TOTP authenticators for use in *nZTA*, see Workflow: Adding TOTP to an Authentication Policy.

Using User Authentication Policies

Users trigger authentication by connecting to the access URL defined in your *authentication policies*. These policies reference *authentication methods*, as described in Creating User Authentication Methods, and cover the primary scenarios for authenticating user connections:

- User enrollment
- User sign-in
- Administrator sign-in

nZTA supports Multi-Factor Authentication for User sign-in and Administrator sign-in type policies only.

nZTA is pre-configured with built-in "default" policies for each of these scenarios.

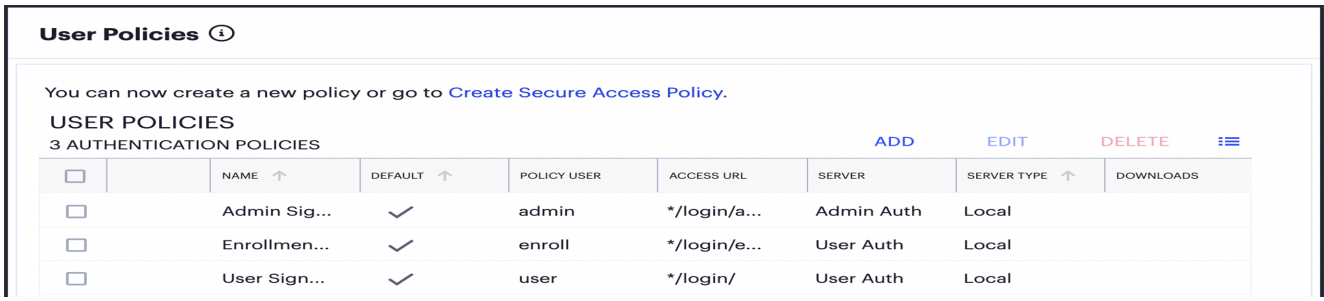

Viewing the default user authentication policies provided by the *Controller*

The access URL defined in each policy is appended to the tenant FQDN and serves as a complete endpoint for that activity. For example, if your tenant FQDN is "example.company.com", the default sign-in endpoint for your end-users is *https://example.company.com/login/* and the default enrollment endpoint is *https://example.company.com/login/enroll/*.

The specific access URL to which you connect for enrollment depends on the *Ivanti Secure Access Client* version installed on the device:

- For *Ivanti Secure Access Client* versions defined as recommended for *nZTA* 21.9.2 or later, or for browser-based enrollment, you enroll a device by connecting directly to the access URL in a **user sign-in** policy. You do not connect to the **enrollment** policy directly (enrollment policies are linked to sign-in policies to handle the case where a connecting device is not yet enrolled).
- For *Ivanti Secure Access Client* versions earlier than those recommended for *nZTA* 21.9.2, you enroll a device by connecting to the access URL in an **enrollment** policy. Thereafter, enrolled devices connect to the access URL in your default **user sign-in** policy.

For details of recommended and supported *Ivanti Secure Access Client* versions, see the *Release Notes*.

The built-in default policies cannot be replaced or deleted, yet can be edited to update the authentication methods used in each case. For standard deployments involving a single authentication workflow for all enterprise users, these policies should be sufficient.

To establish additional authentication endpoints for specific user groups, you can add *custom* authentication policies. To learn more, see Adding Custom Authentication Policies.

Adding Custom Authentication Policies

In addition to the built-in authentication policies, you can add custom policies to enable further group-specific login scenarios. You can add, edit, and delete custom authentication policies as required.

In other words, where you only require a single policy for all enterprise users containing an authentication method suitable for everyone, the built-in default policies are likely suitable for your needs. However, you might want to authenticate certain groups of users against a different authentication method. In this scenario, you can set up custom policies, each containing a separate authentication method suitable for your specific group of users (for example, one policy/method for Sales staff versus another for non-Sales staff).

In a custom authentication policy, you define a unique *Access URL* for the activity (for example, "*/login/saleslogin/" or "*/login/salesenroll/"), and the primary authentication method you want to apply in that case. For Multi-Factor Authentication scenarios, you should also define an appropriate secondary authentication method (based on local or TOTP authentication).

Each time you create a custom authentication policy, you configure the user type (or activity) to which this policy should apply:

- **Enrollment Users**: This policy is intended for enrollment of new end-user devices.
- **Users**: This policy is intended as the sign-in endpoint for enrolled devices.
- **Administrators**: This policy is intended as the authentication endpoint for administrator-level sign-in to the *Controller*.

Multi-Factor Authentication is supported for user sign-in and admin sign-in custom policies only.

Authentication policies defined as a *users* sign-in type contain a link to an *enrollment users* policy that you specify. Through this mechanism, un-enrolled user devices connecting to a sign-in endpoint are instead automatically redirected to the relevant enrollment endpoint. For this reason, make sure you create your enrollment policy BEFORE you create a sign-in policy.

As an example scenario:

1. For enrollment, you create a new custom user policy specifying a *User Type* of "Enrollment Users" and a unique enrollment *Login URL* such as `/login/testenroll`. Then, select your required authentication server.
2. To enable your users to sign in post-enrollment, create a separate custom user policy specifying a *User Type* of "Users". Enter a separate unique *Login URL* representing the sign-in endpoint, such as `/login/testuser`. Then, select your required primary (and optional secondary) authentication server. Finally, link in the *enrollment policy* you created in the previous step.

3. When you add a new *nSA* connection in *Ivanti Secure Access Client*, you specify the sign-in endpoint (`https://<mydomain>.com/login/testuser`). This automatically redirects to `https://<mydomain>.com/login/testenroll` to complete the enrollment process. An enrollment endpoint must not itself be used for device onboarding from *Ivanti Secure Access Client* or through a web browser-based enrollment.

Your custom user sign-in policies can be configured to redirect enrollment to either another custom (enrollment) policy, or to the built-in *Enrollment Signin* policy. The built-in *User Signin* policy is pre-configured with the built-in *Enrollment Signin* policy and cannot be modified.

If you use custom authentication policies, you must make sure these policies are mapped to a valid *Secure Access Policy*. Otherwise, connection attempts to enroll an end-user device to these policy end-points will fail.

For more details on enrolling client devices, see Enrolling a User Device.

To learn more about creating user authentication policies, see Working with User Authentication.

Defining User Authentication Rules and Groups

After you have defined a user authentication service, you can proceed to create **User Rules** and **User Groups**.

A **User Rule** is a set of instructions describing whether *nZTA* should include or exclude a user based on whether the criteria defined in the rule matches the user. For each rule, you specify an attribute of a user's credentials (such as "username") and a value to match against or to avoid. The value can be a wild-card character, a literal string, or a mixture of both. For example, you might specify a rule of "username MATCHING *@example.com" to apply to all users with user names in the example.com domain.

The attribute to match varies depending on whether your rule is to be used against an internally held user list or a SAML-based external authentication service.

Consider specifying separate user rules for each type of user in your organization. Consider also defining rules for each type of application or resource you want to publish. For example, you might require a rule to filter out those users that are not authorized to access a sensitive data application, such as an HR or payroll system, held in a specific datacenter.

A **User Group** contains one or more user rules that, when combined with a user authentication service (see Defining User Authentication), create a complete user authentication policy.

This capability, combined with device policies, can actively fulfill the Users and Devices dimensions of your secure access policy. For more information, see An Overview of Ivanti Neurons for Zero Trust Access.

For logins attempted against a URL endpoint defined in a particular policy, the credentials of the user are checked against the corresponding authentication services and user rules to ascertain whether the user is authorized to access the services published at a *nZTA Gateway*.

To learn more about defining User Rules and Groups, see Creating User Rules and User Groups.

Publishing Applications

Application publishing is central to the configuration of your *nZTA* service. A *nZTA* application can refer to on-premise applications, web pages, or network locations served from your datacenter and cloud infrastructure. *nZTA* can also publish resources based on SaaS applications such as Microsoft O365 and Salesforce.

The process of publishing an application encompasses all four dimensions of secure access: users, devices, applications, and infrastructure. For each application, you define its type and then apply to it user rules, device policies, and *nZTA Gateway* infrastructure governing access to it. For more information on the secure access model, see An Overview of Ivanti Neurons for Zero Trust Access.

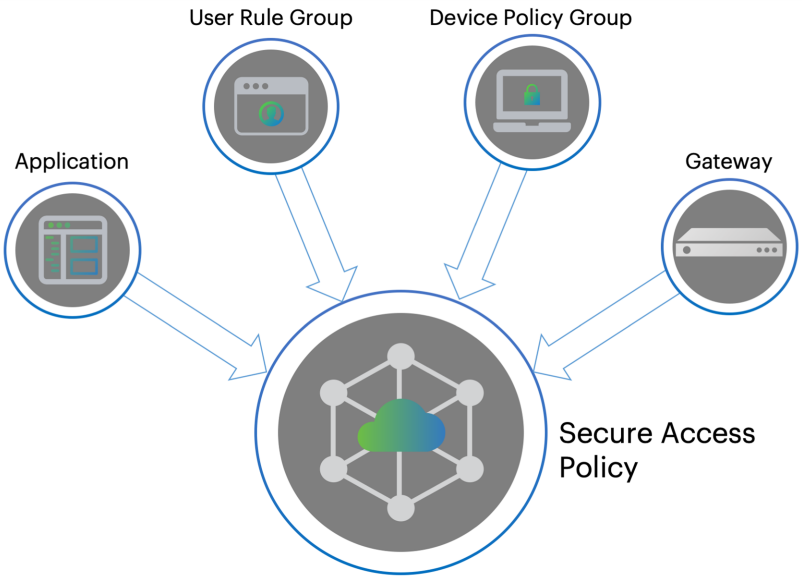

Creating a secure access policy

The Tenant Admin Portal allows you to define a complete *secure access policy* for an application or resource, encompassing the required supporting user rules, device policies, and *nZTA Gateway* infrastructure.

The portal provides a step-by-step workflow to enable you to define and publish your secure access policy in one process. For more information on the **Create Secure Access Policy** workflow, see Summary of Steps to Configure Your nZTA Deployment.

Defining Applications and Application Groups

When you create a new *secure access policy*, you select whether to create a policy for a single application or for a group of related applications. Therefore, the Tenant Admin Portal enables you to define applications individually, and where required to then add one or more application definitions to an *application group*.

An application, or application group, can be associated with only one secure access policy.

When you create a new application definition in the Tenant Admin Portal, make sure you know the following details:

- **The Application Details**: That is, the URI (Uniform Resource Identifier) you use to access the application, based on a Fully Qualified Domain Name (FQDN). *nZTA* supports valid entries that match the following forms:

Application Details

URI Scheme	Example(s)
http://<FQDN or IP>[:<port>][/path]	http://www.example.com, http://www.example.com:80, http://www.example.com/myapp, http://www.example.com:80/myapp, http://192.0.2.0, http://192.0.2.0:80, http://192.0.2.0/myapp, http://192.0.2.0:80/myapp
https://<FQDN or IP>[:<port>][/path]	https://www.example.com, https://www.example.com:443, https://www.example.com/myapp, https://www.example.com:443/myapp, https://192.0.2.0, https://192.0.2.0:443, https://192.0.2.0/myapp, https://192.0.2.0:443/myapp
tcp://<FQDN or IP>[:<port>], tcp://<FQDN or IP>[:<port1,port2>], tcp://<FQDN or IP>[:<portA-portB>], tcp://<FQDN or IP>[:*]	tcp://example.com, tcp://example.com:21, tcp://192.0.2.0, tcp://192.0.2.0:1-65535, tcp://192.0.2.0:80,443, tcp://192.0.2.0:*
udp://<FQDN or IP>[:<port>], udp://<FQDN or IP>[:<port1,port2>], udp://<FQDN or IP>	udp://example.com, udp://example.com:69, udp://192.0.2.0, udp://192.0.2.0:1-65535,

URI Scheme	Example(s)
[:<portA-portB>], udp://<FQDN or IP>[:*]	udp://192.0.2.0:161,162, udp://192.0.2.0:*
ssh://<FQDN or IP>[:<port>]	ssh://example.com, ssh://example.com:22, ssh://192.0.2.0, ssh://192.0.2.0:22
rdp://<FQDN or IP>[:<port>]	rdp://example.com, rdp://example.com:3389, rdp://192.0.2.0, rdp://192.0.2.0:3389
ica://<FQDN or IP>	ica://example.com, ica://192.0.2.0
<FQDN>:<port>	www.example.com:9090
<FQDN>:<port1>,<port2>[,...]	www.example.com:9090,9091,9092
<FQDN>:<portA>-<portB>	www.example.com:9090-9099
<FQDN>:*	www.example.com:*
<IP address block with subnet>	192.0.2.0/8, 192.0.2.0/255.0.0.0
<IP address>[/<subnet>]	192.0.2.1, 192.0.2.1/32, 192.0.2.1/255.255.255.255

IPv6 addresses are not supported when defining application details

- **The access type**. This is one of the following types:
  - *Application-level*: Defines an on-premise or cloud application.
  - *SAML*: Defines a trusted connection between *nZTA* and the application, suitable for SaaS applications.

When using SAML authentication with an application group, make sure you add to a single application group only those applications that use the same SAML authentication source.

For simple application access, you specify the full FQDN, URI or IP address at which the application is accessed. In some cases, you might have multiple applications available at a single primary domain, the full set of which might be unknown to the administrator configuring your *nZTA* service, particularly with SaaS applications. Instead of locating and defining each of these applications individually, you can create a single application definition based on a wildcard-prefixed FQDN (for example, "*.example.com") and instruct *nZTA* to employ *application discovery* on that domain. The method means that a single secure access policy can handle access to all applications *discovered* at the domain.

nZTA records analytics and metrics for all discovered applications. To learn more about how to track usage, see Reviewing Application Usage.

To learn more about the process of creating applications and application groups, see Adding Applications to the Controller.

Defining an On-Premise Application

Before configuring *nZTA* to publish your resource, make sure the resource is accessible from the *nZTA Gateway* you want to use for controlling access. You should obtain the FQDN of the endpoint at which the resource is visible.

nZTA additionally requires a name, description, and icon to represent the resource.

Defining a SaaS Application

The *Controller* uses SAML to provide a secure connection to a SaaS resource. In this scenario, *nZTA* acts as a SAML Identity Provider (IdP), with the resource acting as the SAML Service Provider (SP). To learn more about using SAML, see SAML Authentication.

By selecting SAML as the application access type, the Tenant Admin Portal provides the ability to download the IdP metadata file for uploading to the application SP configuration. This action, in turn, should elicit a SP metadata file for uploading back to *nZTA*.

Defining User Rules and User Rule Groups

A secure access policy links to the same user rule groups as described during the process of configuring user authentication. You create access rules for your enterprise user base and add these rules to a corresponding rule group. You then select a user rule group during the creation of your secure access policy.

For more information, see Defining User Authentication Rules and Groups.

To learn more about the process of defining user rules and rule groups, see Creating User Rules and User Groups.

Creating Device Policies and Device Policy Rules

Device policies define the minimum standard a device must meet to be considered compliant with *nZTA*.

Examples of compliance include:

- All client devices, desktop or mobile, must have anti-virus software installed and enabled.
- A specific application might function securely only on devices using a minimum viable Operating System version.
- Application access on client devices should be restricted to a defined browser type and minimum version.
- Application access should be restricted to a pre-determined range of authorized client device IP addresses.

You create individual device policy through the Tenant Admin Portal and then create / associate device rules to them. As you move on to define your secure access policy, you select a device policy that should be enforced for the selected application, users, and *nZTA Gateway*.

nZTA includes the following built-in device policies:

- For McAfee AntiVirus (Windows/Mac):
    - **McAfeeAVHigh**: This is a high strictness policy designed to ensure devices have anti-virus software installed and active, with a virus signature database not more than 2 days old, and a full system scan having been performed. This policy is attached to the built-in **McAfeeAntiVirusHigh** device rule.
    - **McAfeeAVMedium**: This is a moderate strictness policy designed to ensure devices have anti-virus software installed and active, with a virus signature database not more than 10 days old. This policy is attached to the built-in **McAfeeAntiVirusMedium** device rule.
    - **McAfeeAVLow**: This is a low strictness policy designed to ensure devices have anti-virus software installed and active. This policy is attached to the built-in **McAfeeAntiVirusLow** device rule.

- For Symantec AntiVirus (Windows/Mac):
    - **SymantecAVHigh**: This is a high strictness policy designed to ensure devices have anti-virus software installed and active, with a virus signature database not more than 2 days old, and a full system scan having been performed. This policy is attached to the built-in **SymantecAntiVirusHigh** device rule.
    - **SymantecAVMedium**: This is a moderate strictness policy designed to ensure devices have anti-virus software installed and active, with a virus signature database not more than 10 days old. This policy is attached to the built-in **SymantecAntiVirusMedium** device rule.
    - **SymantecAVLow**: This is a low strictness policy designed to ensure devices have anti-virus software installed and active. This policy is attached to the built-in **SymantecAntiVirusLow** device rule.
- For TrendMicro AntiVirus (Windows/Mac):
    - **TrendMicroAVHigh**: This is a high strictness policy rule designed to ensure devices have anti-virus software installed and active, with a virus signature database not more than 2 days old, and a full system scan having been performed. This policy is attached to the built-in **TrendMicroAntiVirusHigh** device rule.
    - **TrendMicroAVMedium**: This is a moderate strictness policy rule designed to ensure devices have anti-virus software installed and active, with a virus signature database not more than 10 days old. This policy is attached to the built-in **TrendMicroAntiVirusMedium** device rule.
    - **TrendMicroAVLow**: This is a low strictness policy rule designed to ensure devices have anti-virus software installed and active. This policy is attached to the built-in **TrendMicroAntiVirusLow** device rule.
- For device modification testing:
    - **AndroidRootRulePolicy**: This policy is designed to ensure that an Android device has not been rooted. This policy is attached to the built-in **AndroidRootRule** device rule.
    - **IOSJailBreakRulePolicy**: This policy is designed to ensure that an iOS device is not subject to a jailbreak. This policy is attached to the built-in **IOSJailBreakRule** device rule.

- For risk-based vulnerability:
    - **RiskSenseCriticalNotify**: This policy adds robust risk-based vulnerability prioritization and remediation capabilities to Ivanti Neurons for Patch Intelligence. This policy is attached to the built-in **RiskSenseCriticalNotifyRule** device rule.

You can create additional device policies to match the applications used by your organization. You then associate device policies with your secure access policies as required.

Policies can be based on the following types:

Policy Type	Description
Antispyware	Checks compliance to designated anti-spyware requirements.
Antivirus	Checks compliance to designated anti-virus requirements.
CVE check	Checks for protection against a list of publicly disclosed Common Vulnerability and Exposure (CVE) notices (Windows client devices only).
Command	Runs a command on the client device to check against an expected value (macOS client devices only).
File	Checks for the existence of a known file on the client.
Firewall	Checks compliance to designated firewall requirements.
Location	Checks the client device's geographic location matches, or avoids, a list of defined locations.
Hard Disk Encryption	If encryption software is installed on the client device, this rule type checks the device's hard disks for applied encryption.
Mac Address	Checks the client device's MAC address.
Netbios	Checks the client device's Netbios domain name.
Network	Checks the client device complies with a defined IP address and netmask range.
OS	Checks the client device’s Operating System meets a defined minimum standard.
Process	Checks for the existence of a known process on the client.
Port	Checks the client device's network interface ports.
Patch Management	If patch management software is installed on a client device, this rule type checks for the existence of missing software patches.

Policy Type	Description
Registry	Checks for a value in a registry key (Windows client devices only).
Risk Sense	Supports Allow access, Block access and Notify based on the risk level.
Time of day	Checks resource access requests against compliance with a time-based access schedule.
System Integrity	Checks the system integrity of the client device (macOS client devices only).

Enrolling a User Device

To use *nZTA*-controlled resources, end users must first enroll their devices with the *Controller*. This process installs or updates the *Ivanti Secure Access Client* software and establishes a connection to the *Controller* in order to obtain policies and details for a user's authorized resources. *Ivanti Secure Access Client* uses this configuration to establish a secure connection to the *nZTA Gateway*s you deploy to control access to your applications. Through this process, the user is provided a seamless connection to the resources they need and is never aware of the location or extent of the organization's application infrastructure.

For security reasons, only the authorized user account used to enroll a device is subsequently permitted to sign-in to *nZTA* on that device.

A new user might arrive at this scenario from one of the following routes:

- An existing Pulse Secure or *Ivanti* user, with a previous *Ivanti Secure Access Client* connection to Ivanti Connect Secure (ICS) or similar, see Existing ICS Users.
- A first time *nZTA* user, with no previous Pulse Secure or *Ivanti* software installed, see First Time Users.
- An existing *nZTA* user enrolling a new device, or upgrading a previous version of *Ivanti Secure Access Client*, see Existing nZTA Users.

To learn more about the processes of enrolling supported mobile and desktop clients, and to see how an admin can manage the list of enrolled devices, see Enrolling Mobile/Desktop Clients.

To learn more about using mobile and desktop clients with *nZTA*, see Using Ivanti Secure Access Client with nZTA and Enrolling Mobile/Desktop Clients.

Existing ICS Users

Your users might have a previous version of *Ivanti Secure Access Client* installed if, for example, they are existing ICS users.

To enroll existing ICS users into *nZTA*, a ICS administrator must first push out the *nZTA*-ready edition of the client software to the user base. An admin uploads the new client software to the ICS server and activates the *nZTA*-ready version of *Ivanti Secure Access Client* from the ICS management console in the same way as any other version. This process ensures that when your users next activate a *Ivanti Secure Access Client* connection to the server, their device is prompted to download and install the new version.

For more details on this process, see the Ivanti Connect Secure documentation at https://help.ivanti.com.

After the new *nZTA*-ready version of *Ivanti Secure Access Client* is installed, the user can configure a *nZTA* connection using the same process used for other, existing, connections. To create a *nZTA* connection, compatible *Ivanti Secure Access Client* versions offer a specific connection type: "**Zero Trust Access**".

The tenant admin must then supply the *access URL* specified in the relevant user sign-in authentication policy to their users to create the new *nZTA* connection. To learn more about authentication policies, see Using User Authentication Policies.

First Time Users

When enrolling a new device on which no previous *Ivanti Secure Access Client* is installed, an authorized user contacts the *Controller* through the access URL of a *user sign-in authentication policy* to activate an initial first-time enrollment of their client device. The *Controller* responds to a valid enrollment request by activating a download of *Ivanti Secure Access Client* along with a suitable client certificate.

After *Ivanti Secure Access Client* is installed, a secure connection request is attempted with the *Controller*. The request is validated against the designated authentication policy applicable to that combination of user and device and, where successful, a connection profile is downloaded to the client. This profile enables *Ivanti Secure Access Client* to set up a secure tunnel directly to the *nZTA Gateway* serving the resource set the client is authorized to view.

Existing *nZTA* Users

Existing *nZTA* users enrolling a new device follow the same procedures described for a first time user, but based on the additional conditions noted here.

The access URL to which you connect for enrollment depends on the *Ivanti Secure Access Client* version installed on the device:

- For *Ivanti Secure Access Client* versions defined as recommended for *nZTA* 21.9.2 or later, or for browser-based enrollment, you enroll a device by connecting directly to the access URL in a **user sign-in** policy. You do not connect to the **enrollment** policy directly (enrollment policies are linked to sign-in policies to handle the case where a connecting device is not yet enrolled).
- For *Ivanti Secure Access Client* versions earlier than those recommended for *nZTA* 21.9.2, you enroll a device by connecting to the access URL in an **enrollment** policy. Thereafter, enrolled devices connect to the access URL in your default **user sign-in** policy.

For details of recommended and supported *Ivanti Secure Access Client* versions, see the *nSA Controller Release Notes*.

After you have enrolled the new device, *Ivanti Secure Access Client* is installed and configured with the policies and settings relevant to the device type. Your application and resource access rights should be duplicated to the new device.

If a user device is currently using a Beta version of the *nZTA*-ready *Ivanti Secure Access Client*, *Ivanti* advises to remove the *nZTA* connection from *Ivanti Secure Access Client* and to re-perform the enrollment procedure through a web browser. For more details, see your support representative.

Viewing Licensing/Subscription Usage

Licenses and subscriptions are added to the *Controller* by *Ivanti*.

The **Subscriptions** page displays your licenses and subscriptions that are active on the *Controller*. To access this page, select **Administration > Subscriptions**.

The **Summary** tab displays for each subscription or license:

- License/Subscription high-level details, including dates and usage metrics.
- One or more descriptions of the features in the license or subscription. Where there are multiple features, use **Previous** and **Next** to navigate.

When any defined limit on the license or subscription is met, a message appears:

- At 75% utilization of seats, an information message appears at the bottom of the screen. You can optionally select **Close**.
- At 90% utilization of seats, a modal message appears at login. Select **Dismiss** to clear the message.
- When 25% of the duration of the license or subscription remains, a modal message appears at login. Select **Dismiss** to clear the message.

Named Users lists users and their devices registered on the *Controller*.

Summary information for *Controller* licenses and subscriptions is displayed at the top of the page:

- The total number of seats from all licenses or subscriptions.
- The number of named users on the *Controller*. Each of these is listed in the table below the summary.
- The percentage of seats consumed.

For each named user, the following information is displayed:

- The name of the user.
- The number of devices enrolled for that user.
- The elapsed time since the most recent session on a device enrolled for that user. If this is greater than 3 weeks, this is displayed in orange.

- The timestamp for the start of the most recent login session.

Summary of Steps to Configure Your *nZTA* Deployment

The remainder of this guide describes each of the processes you need to follow to fully configure your *nZTA* deployment.

In summary, the main steps are as follows:

1. Log in to the *nZTA* Tenant Admin portal
2. Create and deploy your *nZTA Gateway*s
3. Define your user authentication methods and policies
4. Define your user rules and user groups
5. Define your device policies and policy groups
6. Define your applications and application groups
7. Create your complete secure access policy.

The remainder of this guide describes each of these steps in detail.

This guide also provides guidance on client device enrollment (see "Enrolling Ivanti Secure Access Client" on page 508), usage (see Using Ivanti Secure Access Client with nZTA), and a description of the tools you can use to monitor your *nZTA* services (see Using the Insights Menu to Monitor User Activity and Service Usage).

Logging in as a Tenant Administrator

Preparing to Login

As a Tenant Admin, you can configure *Ivanti Neurons for Zero Trust Access* (*nZTA*) to support the Gateways, users, devices, policies and resources that are required for your organization's *nZTA* implementation.

To log into the *Controller*, you require a Tenant Admin login.

All Tenant Admin accounts are set up by the *Ivanti* DevOps team. Once your Tenant Admin account has been created, you will receive an email which describes how to log into the *Controller* as a Tenant Admin.

You can then proceed to login to the *Controller*, see Logging into the Controller as a Tenant Admin.

Logging into the *Controller* as a Tenant Admin

Before you can log in as a Tenant Admin, you will receive an email from the *Ivanti* DevOps team. This email contains:

- Your Tenant Admin user name.
- Your password.
- The *nZTA* domain. That is, the FQDN of the *Controller*.
- A hyperlink to start the login process.

To log into your Tenant Admin account:

1. Click the hyperlink in your email.

   The administrator login page appears.

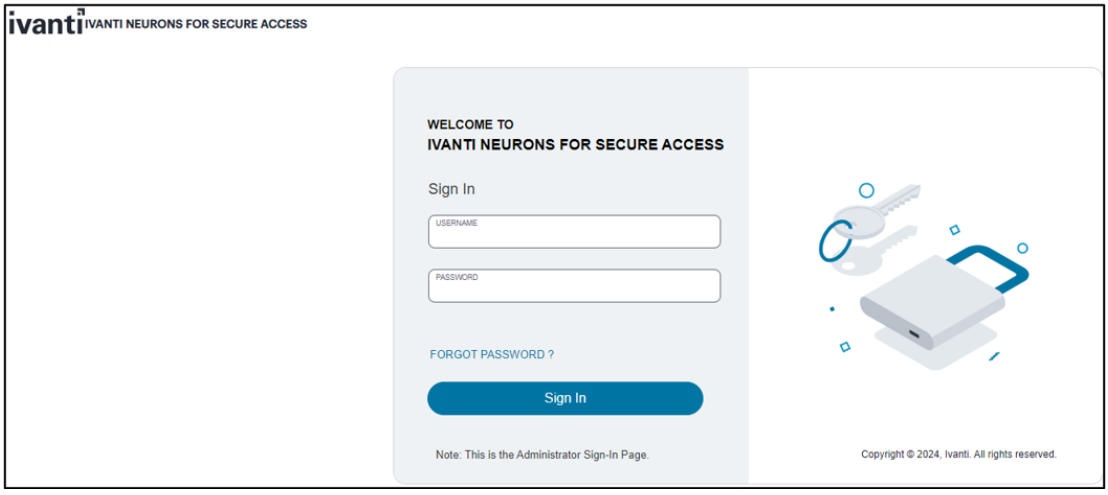

Tenant Admin Login Page

2. Log in using your supplied Tenant Admin credentials.

   The following default timeouts are used for all Admin sessions:

   - The idle timeout is 10 minutes.
   - The session timeout is 60 minutes.

   To configure Admin session timeout values, see Configuring Session Timeouts.

3. If *nZTA* requests it, specify a new password for your account.

   Once this procedure is complete, you access the *nZTA* graphical interface as an admin user.

The graphical user elements that appear depend on your configured state.

- When you log in for the first time, and until *nZTA* is minimally configured, a **Welcome** dialog appears, which leads to the **Secure Access Setup** (Onboarding) wizard. See Working with the Onboarding Wizard.
- When you log into a configured system, the *nZTA* **Network Overview** page appears, see Viewing the nZTA Network Overview.

To reset a forgotten password, click **FORGOT PASSWORD**. This link presents a credentials form through which you enter a Username and Email address. If the entered credentials match a registered administrator account, *nZTA* emails a password reset link to the entered address allowing the recipient to create a new password.

Working with the Onboarding Wizard

When you log in for the first time, and until *nZTA* is a configured system, a **Welcome** dialog appears. To proceed, select **Configure Now**.

The first step enables you to configure a custom domain for this subscription:

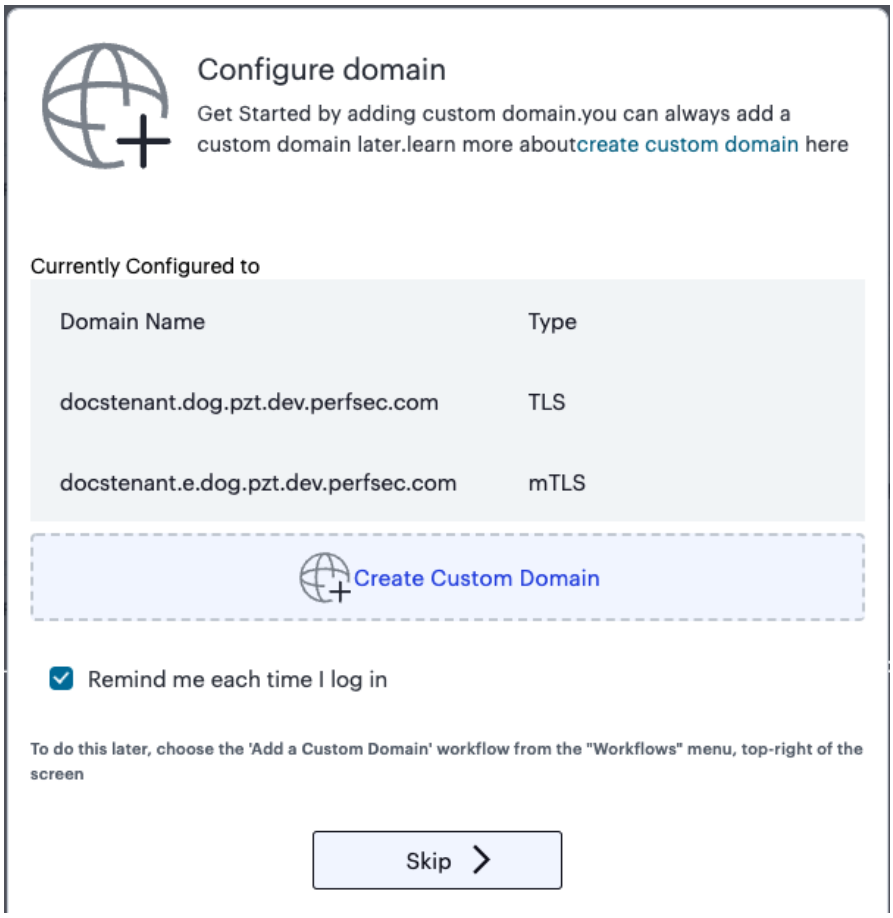

Onboarding Wizard - configuring a custom domain

The current domain (in both TLS and mTLS form) is displayed, along with the option to configure a custom domain. To skip this step and continue using the default domain, select **Skip**.

To configure a custom domain, select **Create Custom Domain** to start the *Create Custom Domain* workflow. For more details on this workflow, see Specifying a Custom Domain.

After you have configured a custom domain, or if you chose to skip ahead, the **Secure Access Setup** (Onboarding) wizard appears.

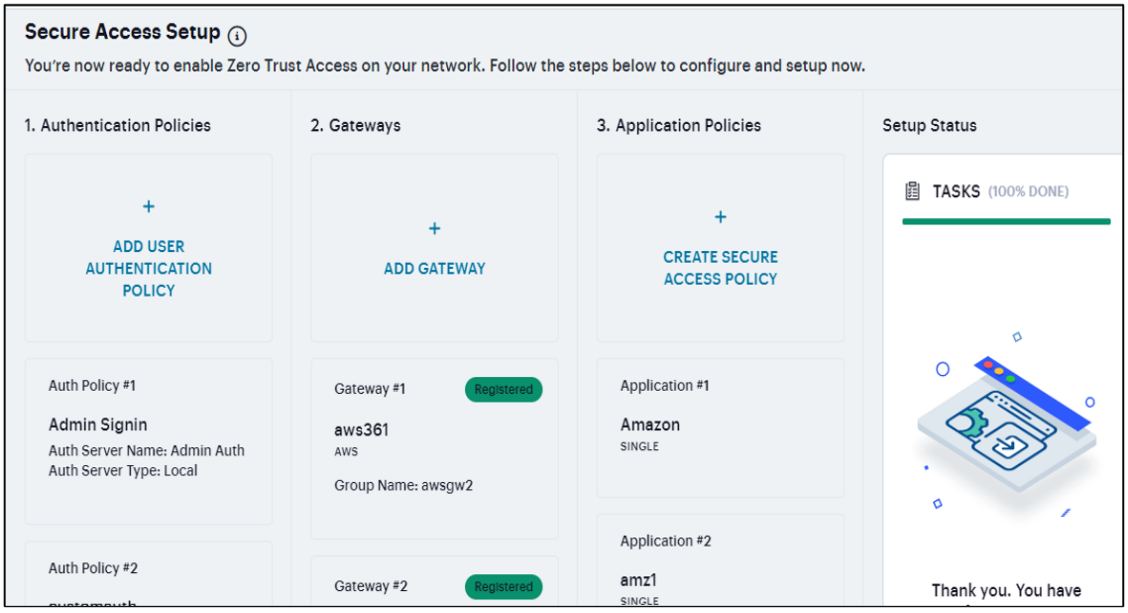

The Secure Access Setup (Onboarding) Wizard

This wizard enables you to configure the required elements of *nZTA* using a number of pages and workflows:

- **Add User Authentication Policy**. This displays the **User Policies** page.

  Local authentication policies are present by default, which can be used immediately.

  If you choose to use the default local authentication policies, you can proceed directly to the **Add Gateway** step.

  If you choose to create your own local authentication policies, or to immediately implement SAML authentication, these must be performed separately from the **Onboarding** wizard, see Working with User Authentication.

- **Add Gateway**. This displays the Gateway Network Configuration workflow, see Working with Gateways.

- **Application Policies**. This displays the Create Secure Access Policy workflow, see Creating Device Policies and Device Policy Rules.

As you complete each steps, the **Setup Status** indicates the percentage of **Tasks** that are complete.

After all tasks are complete, click **Go to Dashboard**.

You can also start the Onboarding wizard from the **Secure Access > Onboarding** menu option.

Viewing the *nZTA* Network Overview

After you log in to the Tenant Admin Portal following successful completion of the *Onboarding Wizard*, *nZTA* displays the **Network Overview** page. This serves as the *home page* for your portal, and provides an overview of user and service activity across your enterprise.

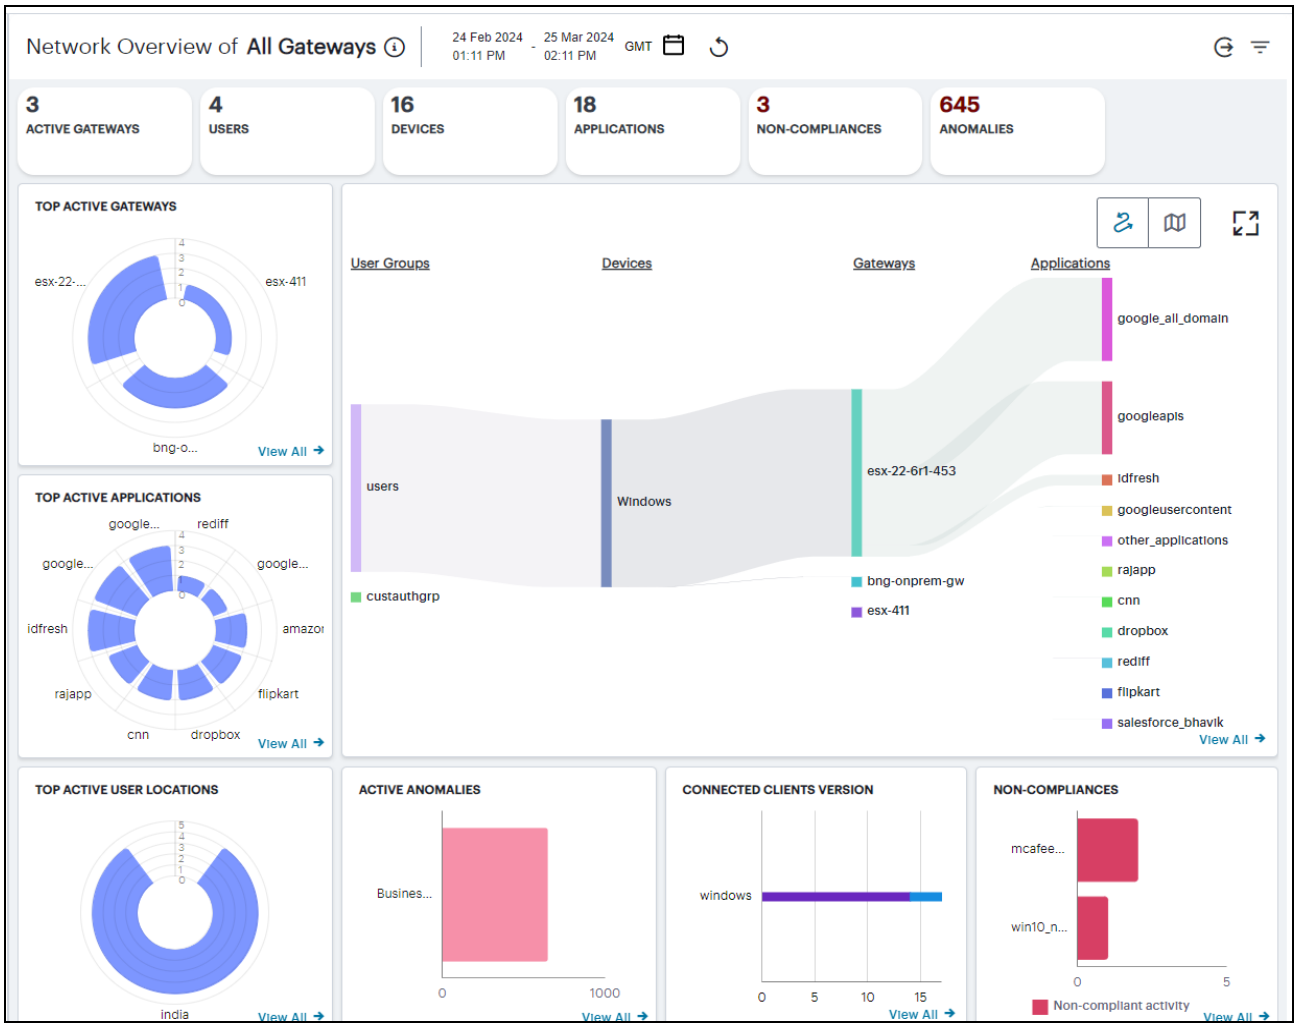

The Network Overview Page

To return to this page any time, click the **Insights** menu icon in the *nZTA* menu and select **Overview**. Alternatively, click the banner at the top.

From this page, you can view and configure all functions and capabilities allowed through your subscription and role. Using the *nZTA* menu at the left-hand side, choose from:

- The **Show/Hide** menu icon, providing the ability to show or collapse the *nZTA* menu tree:

  

  Showing or hiding the *nZTA* menu system

- The **Insights** menu icon, providing access to the analytics and monitoring components of the *nZTA* portal:

  

  Accessing the Insight menu

  To learn more about the functionality offered by this menu, see Using the Insights Menu to Monitor User Activity and Service Usage.

- The **Secure Access** menu icon, providing access to configure the individual components that comprise your *Secure Access Policies*:

  

  Accessing the Secure Access menu

- The **Administration** menu icon, providing access to administrative functions related to your *nZTA* subscription:

  

  Accessing the Administration menu

The chapters in this guide cover each of these functions in detail.

Changing the UI Theme

nZTA offers two themes for your UI display:

- Light theme:

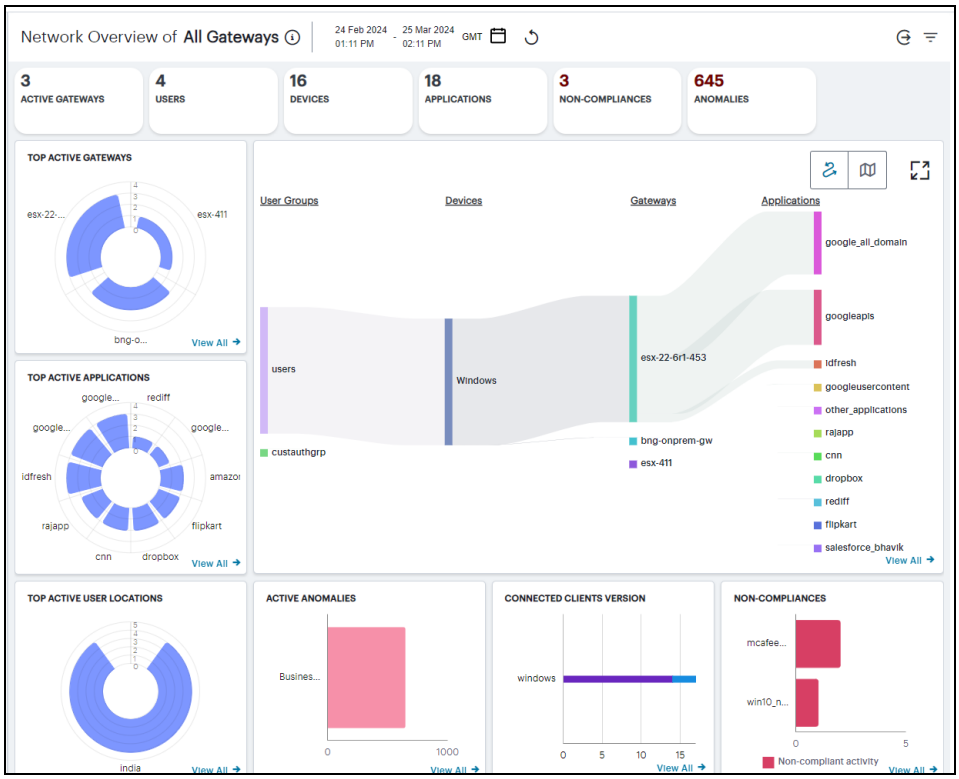

The Network Overview Page - light theme

- Dark theme:

The Network Overview Page - dark theme

To change the current theme, which remains in place through subsequent logins, use the **Settings** menu:

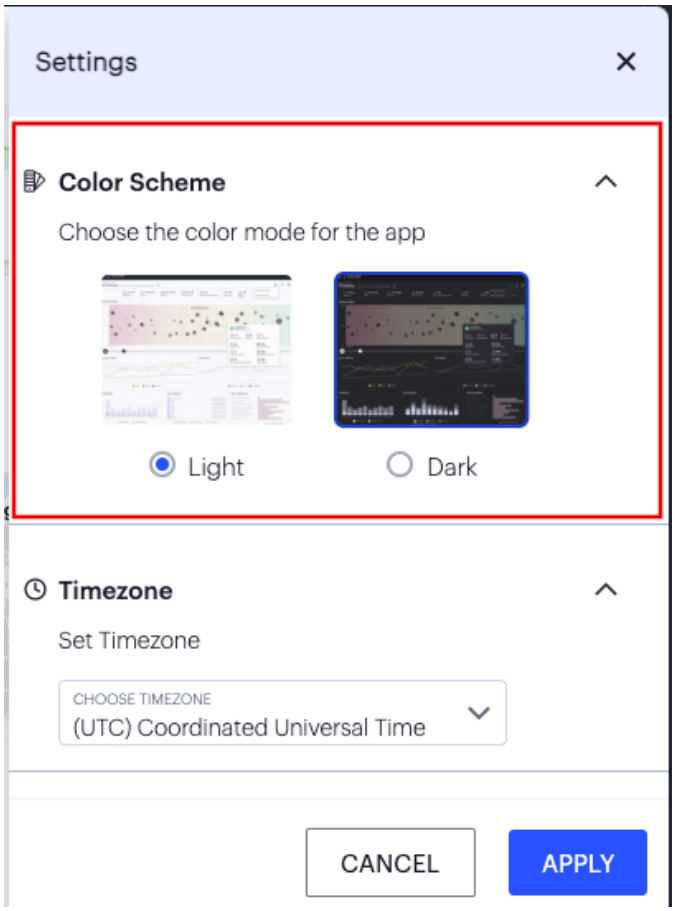

Changing the UI theme

Through the **Color Scheme** panel (indicated), click **Light** or **Dark** to switch between themes.

Setting the Timezone

To configure the default timezone for this admin login account, use the **Settings** menu:

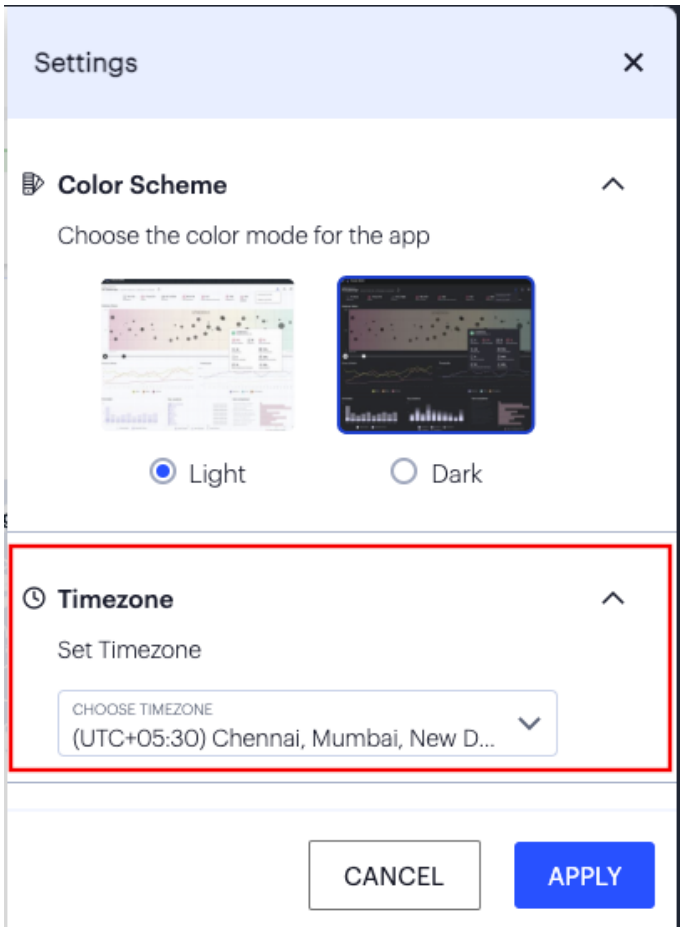

Configuring the default timezone

Choose a timezone in the provided drop-down selector, then click **Apply**.

The configured timezone affects the display of data on all **Insights** pages, and each admin login account within a tenant deployment has their own specific timezone configuration. Changes to the timezone persist across login sessions, and the default setting is *UTC (Coordinated Universal Time)*.

The current timezone can be observed through the date-time selector at the top of each **Insights** page.

Changing the timezone can affect the data displayed in each chart or graph. For example, a recently-observed non-compliance event involving a client device in the GMT timezone that appears in the *Last Hour* view (when using GMT (UTC + 00:00) as your configured timezone) might then only appear in the *Last X Hours* view when you switch your timezone to IST (UTC+05:30).

Configuring Session Timeouts

To configure timeout values for admin and user sessions, use the **Settings** menu:

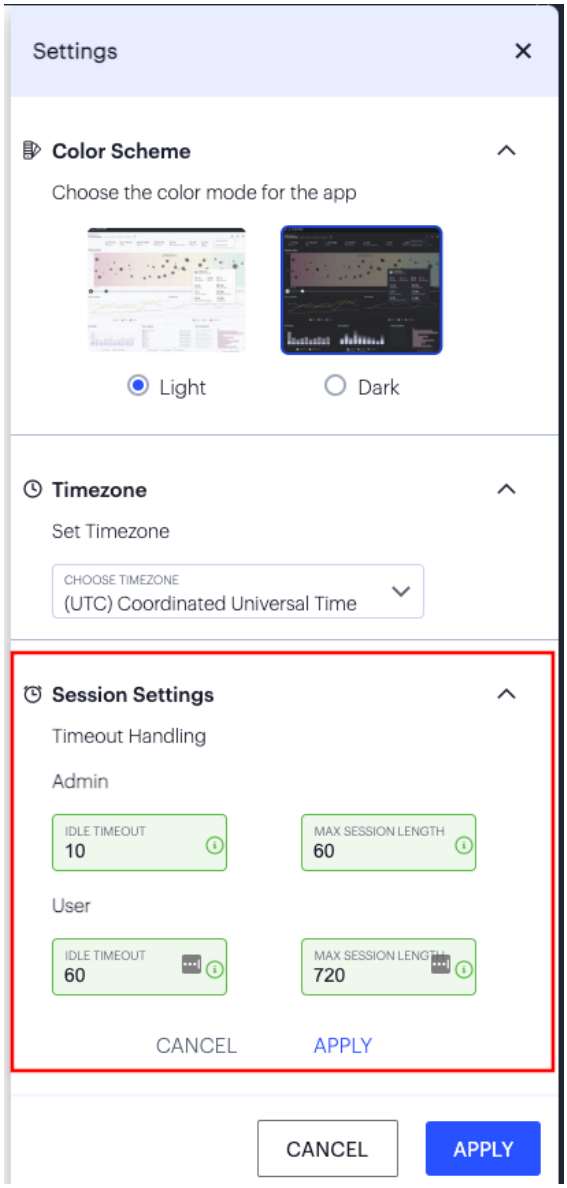

Configuring timeout values for admin and user sessions

Through the **Session Settings** panel (indicated), you can set the following timeout values:

- **Admin Idle Timeout**: the time, in minutes, after which the admin login session to the Tenant Admin Portal times out due to inactivity. (default: 10)

- **Admin Max Session Length**: the time, in minutes, after which the admin login session to the Tenant Admin Portal ends and must be re-authenticated. (default: 60)
- **User Idle Timeout**: the time, in minutes, after which the user login session to *nZTA* times out due to inactivity. (default: 60)
- **User Max Session Length**: the time, in minutes, after which the login session to *nZTA* ends and must be re-authenticated. (default: 720)

To apply your changes, click **APPLY**.

To use these settings, your configured *nZTA Gateway*s must all meet minimum version requirements for session control. *nZTA* disables the panel and displays a warning message if this is not the case.

To learn more about user sessions and the effect of configured timeouts, see Using Ivanti Secure Access Client with nZTA.

Resetting All Filters and Selections

Each page in the *Insights* menu allows data filtering, enabling you to observe and monitor only the analytics and log data you want. Each filter or selection feature includes its own *reset to default value* option. To learn more about the *Insights* menu and the analytics capabilities of *nZTA*, see Using the Insights Menu to Monitor User Activity and Service Usage.

Should you want to reset all filters and data selection criteria across the Tenant Admin Portal in one go, use the **Settings** menu:

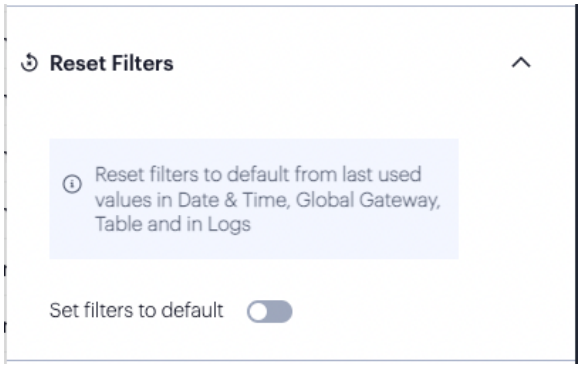

Resetting all filters and data selection criteria across the Tenant Admin Portal

Through the **Filter Reset** panel, you can remove all filters and data selection criteria immediately, returning each page to its default view. Each admin user stores filters and selections separately, so this function affects the current user only.

To activate, enable **Set Filters to default**, then select **APPLY**.

Logging out of the *Controller*

To log out of the *nZTA* Tenant Admin Portal and end the current session, click the *Profile* icon and select **Logout**.

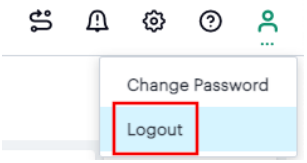

Configuring CASB/SWG

Proxy auto-configuration (PAC) file instructs a browser to forward traffic to a proxy server, instead of directly to the destination server. You can obtain the PAC URL for the Lookout tenant by clicking the PAC file icon from the Secure Cloud Workspace menu.

To get the PAC file URL from Lookout Tenant:

1. From the nZTA menu, select **Integrations** > **CASB/SWG**.

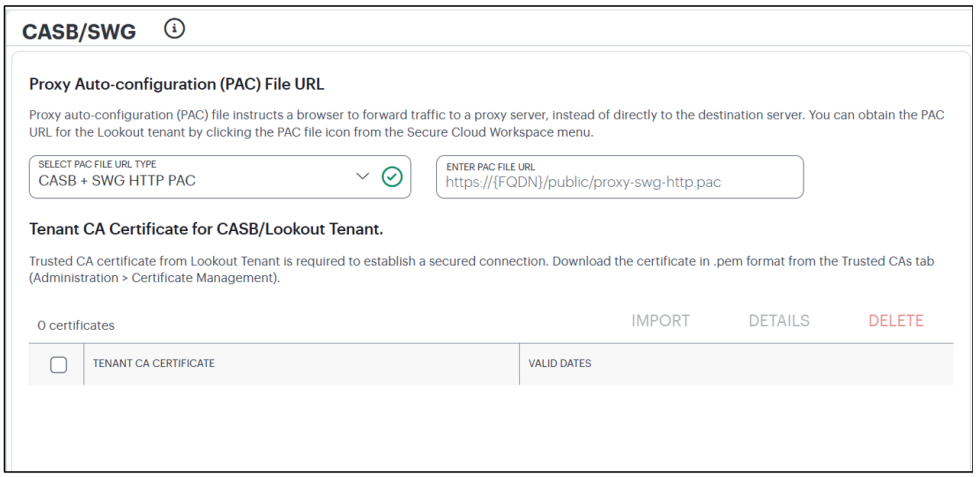

   Configuring CASB/SWG

2. Click the PAC file icon located at the upper-right corner of the page.
3. Hover over the PAC file to apply to your system configuration and click its clipboard copy icon to copy PAC URL.
4. Select PAC file URL type and paste the copied PAC URL in the "Enter PAC file URL" field.
5. Click **Save**.

HTTPS PAC is recommended over HTTP PAC because the tenant identification process is transparent to the end user, and it does not require additional input from the end user.

Trusted CA certificate from Lookout Tenant is required to establish a secured connection. Download the certificate in the .pem format from the Trusted CAs tab (Administration > Certificate Management).

To get the Trusted CA cert from Lookout Tenant and import to the user device:

1. In the Administration > Certificate Management page, click the **Trusted CAs** tab.
2. Download the certificate and save it in the .pem format.
3. In the nZTA tenant SWG/CASB page, click **Import** to import the .pem file.
4. Click **Save**.

For more details about Lookout SWG/CASB Forward Proxy integration, refer https://help.ivanti.com/ps/help/en_US/nSA/22.x/lookout-dep/landingpage.htm.

Integrating Ivanti Neurons for MDM with nZTA

Introduction

Ivanti Neurons for MDM enables to securely access and protect data. It validates the device to ensure that only authorized users, devices, apps, and services can access business resources.

Ivanti NMDM is an independent entity which has a tenant running on the cloud and the same will be used to communicate to the nZTA by adding the details of the NMDM in the nZTA, see "Integrating Ivanti NMDM with nZTA" on page 78.

Neurons for MDM integration with Neurons for ZTA achieves the below goals:

- Supports compliance check for mobile devices during login and application access.

This section explains steps involved in setting up NMDM cloud and configuring NMDM in nZTA.

Prerequisites

- "Configuring Ivanti NMDM Cloud" below

Configuring Ivanti NMDM Cloud

As a prerequisite, an admin needs to configure Ivanti NMDM cloud configuration.

To configure Ivanti NMDM cloud:

1. Log in to Ivanti NMDM server.
2. In the Ivanti Neurons for MDM menu, select **Apps > App Catalog**.

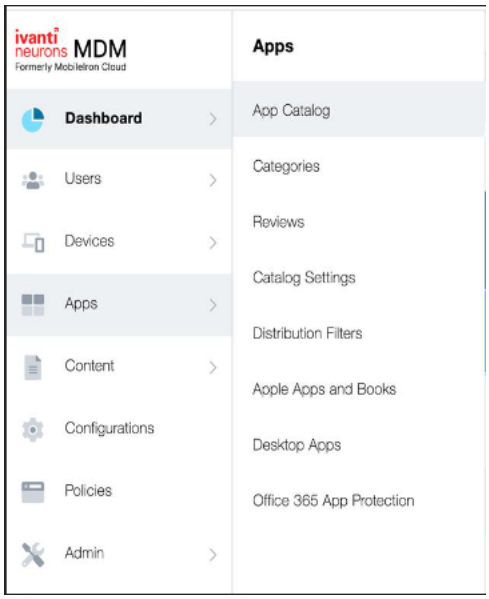

3. Search for *Ivanti Secure Access Client iOS* app. Select the app and click **Next**.

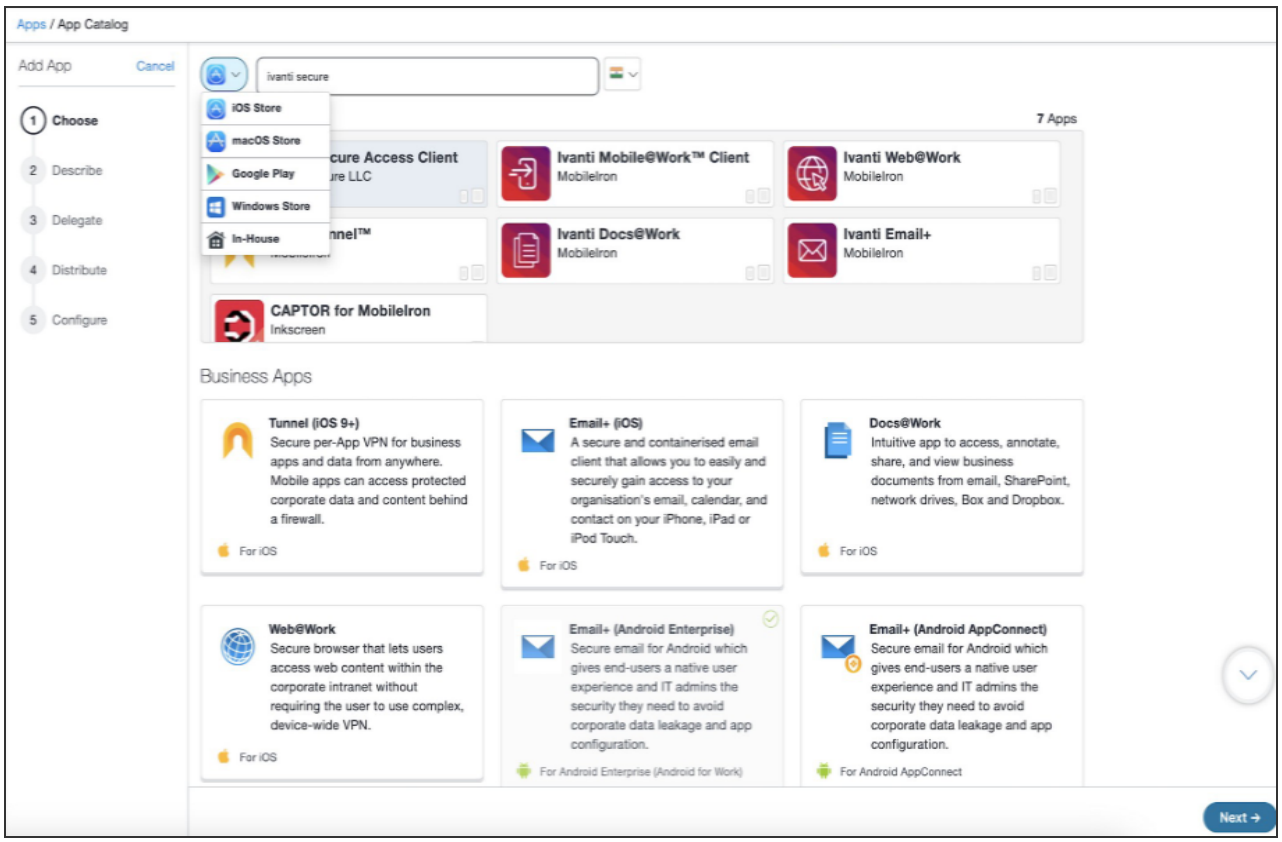

4. In the App Information page, specify **Launch URL** and click **Next**.

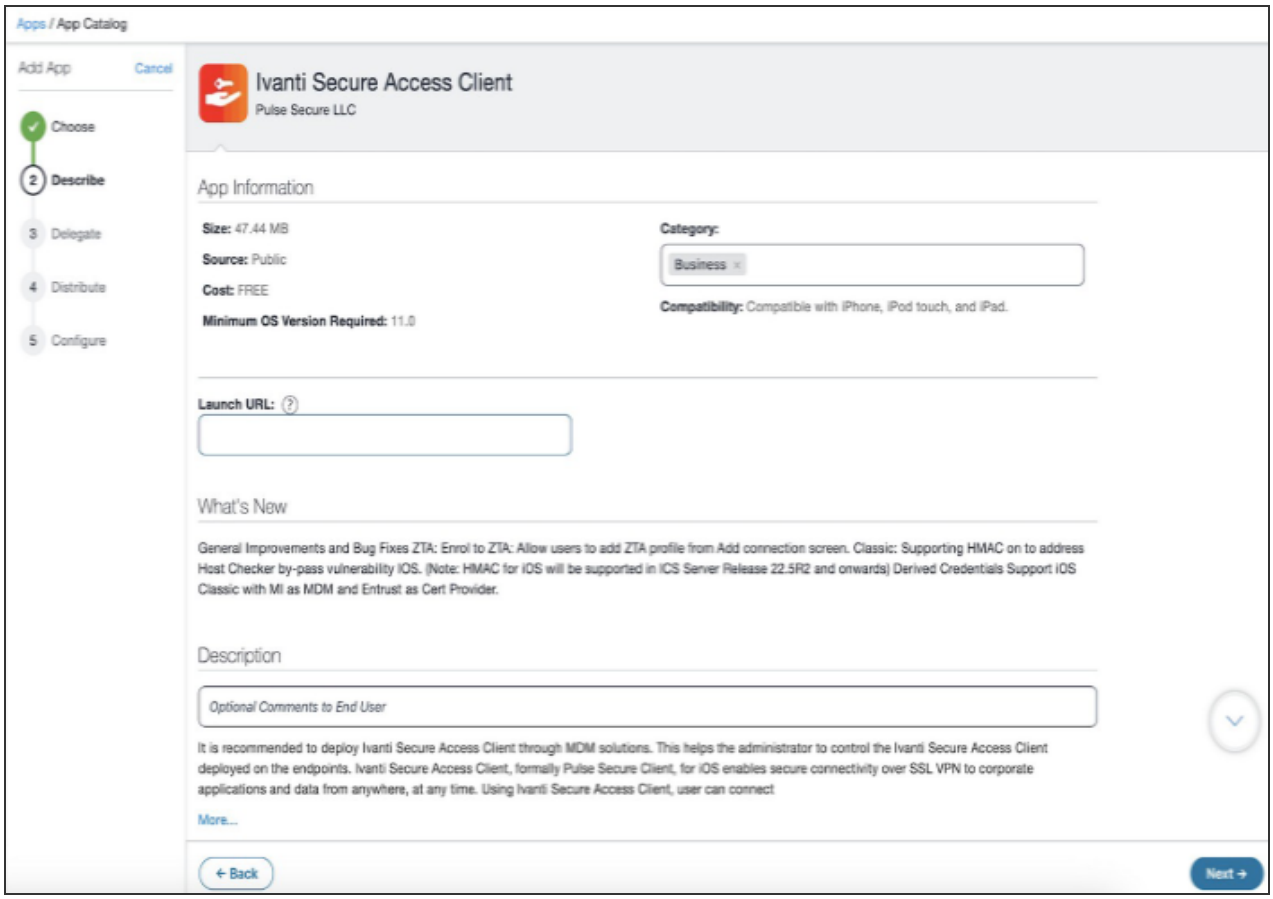

5. In the App Delegation page, select the **Do not delegate the app** option and click **Next**.

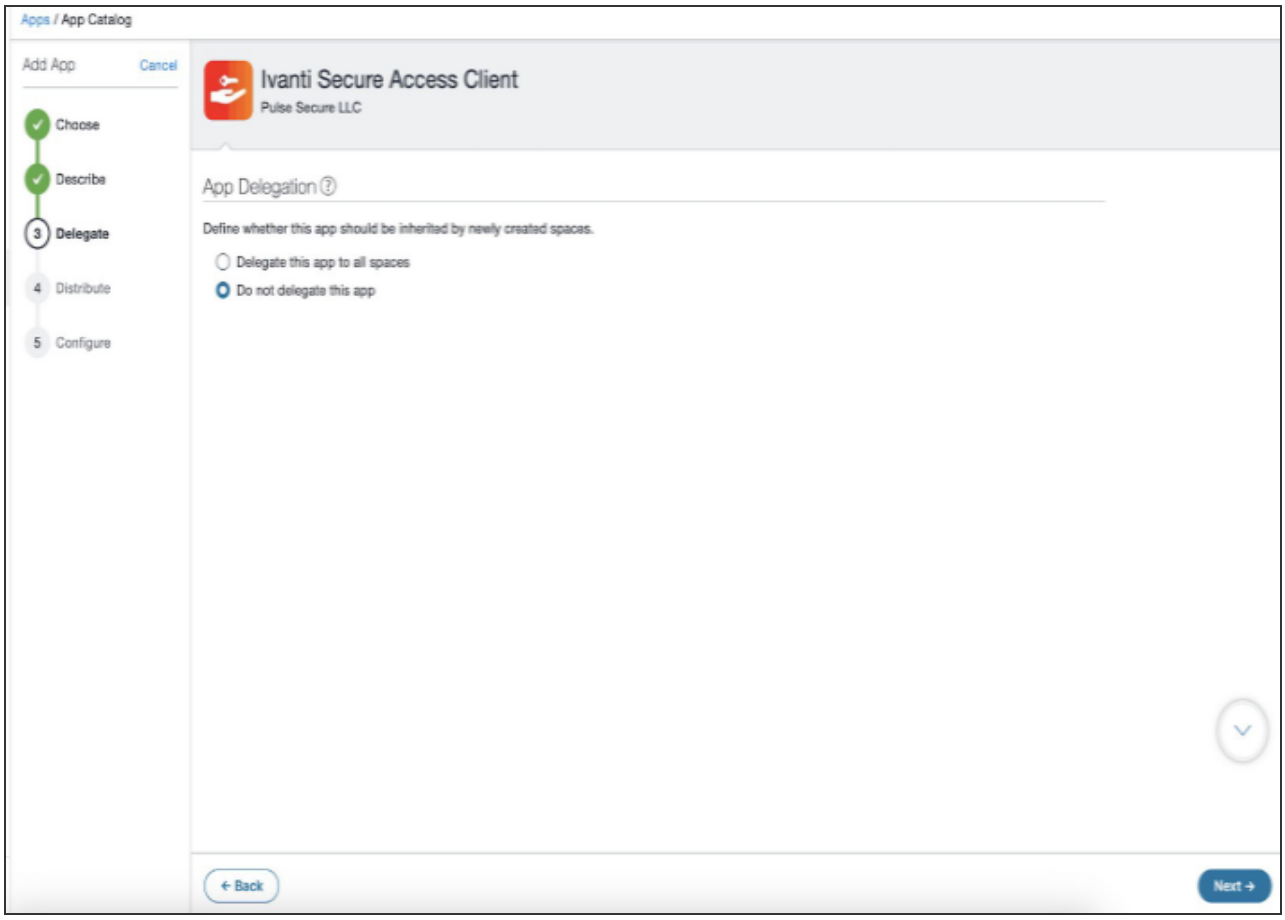

6. Choose a distribution level for this configuration of the app and click **Next**.
    - To Everyone with App - The app is added to all the user compatible devices.
    - To No One - The app is staged for distribution at a later date.
    - Custom distribution - Select one of the options from "User/User Groups" or "Device/Device Groups".

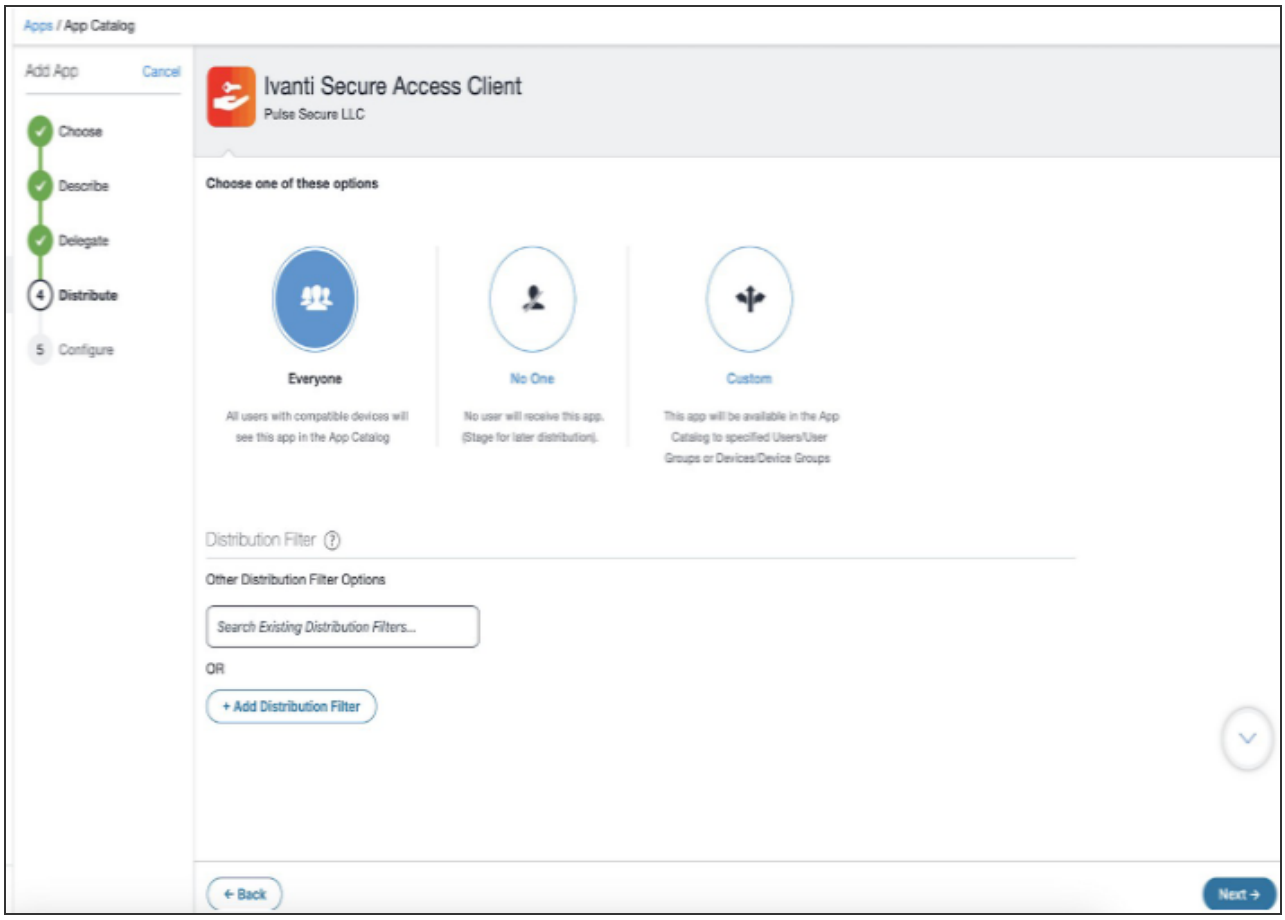

7. In the Configuration Setup page, create a new configuration using plistfile/device attribute ${deviceGUID} of type string. Click **Next**.

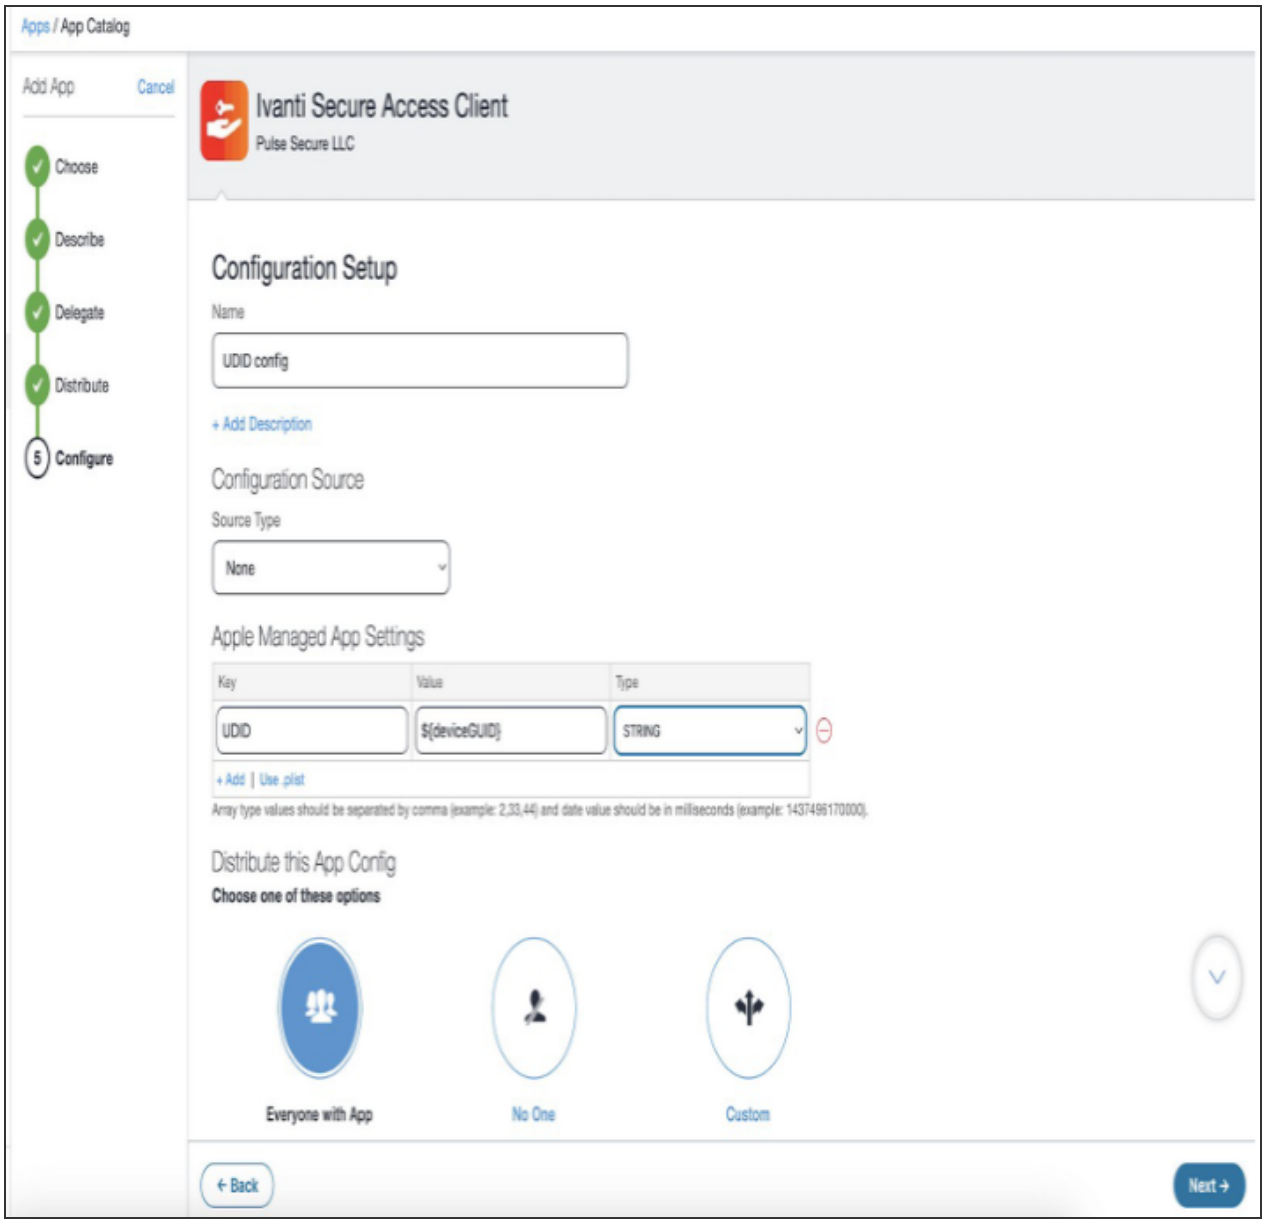

8. In the App Configuration page, select **Apple Application Management Configuration Settings** and click **Done**.

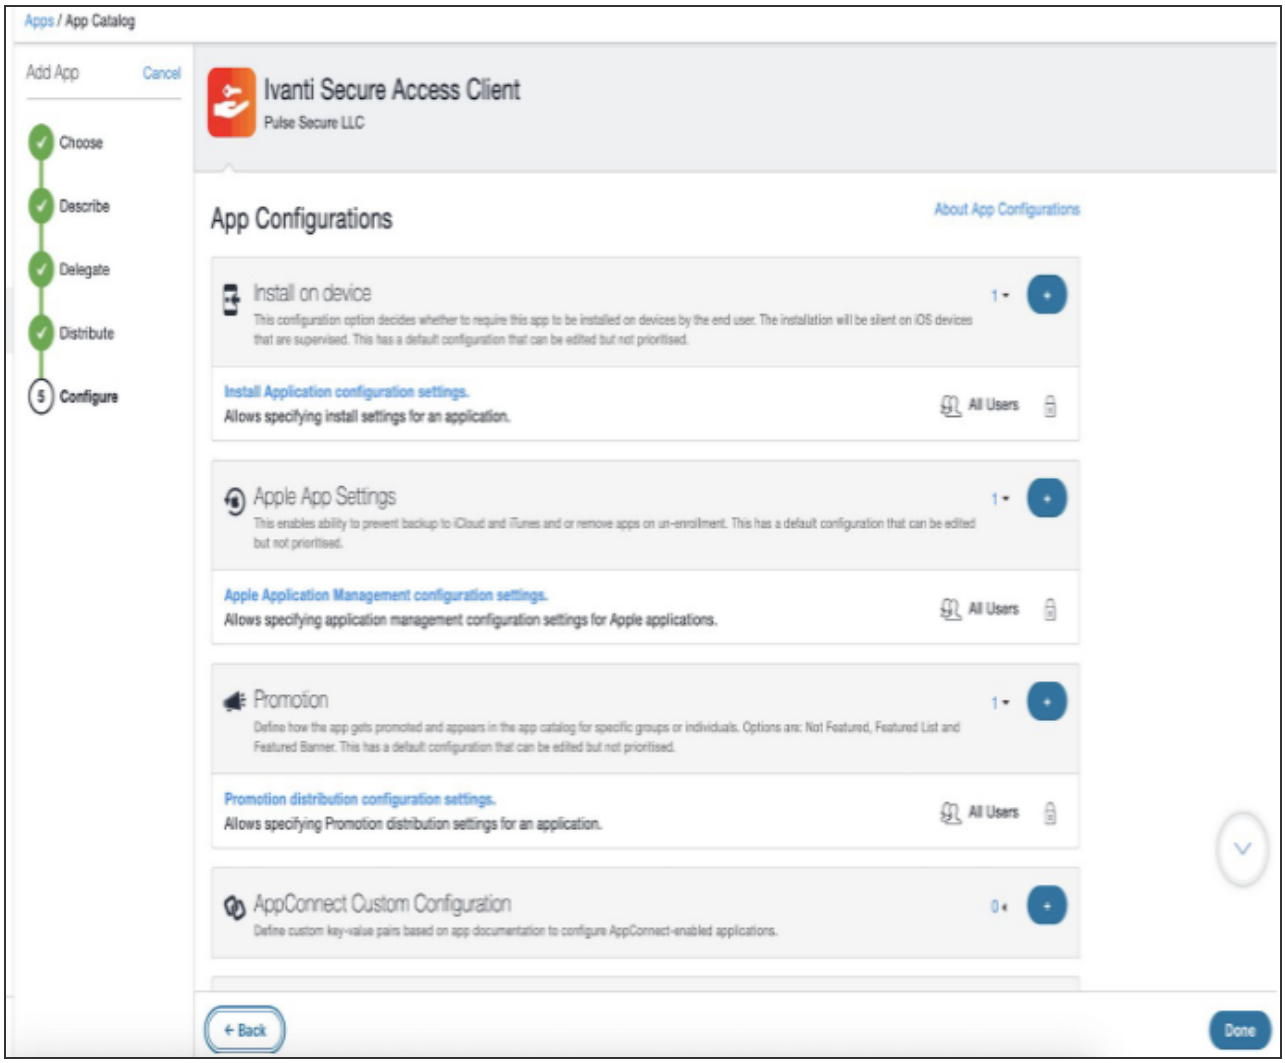

The app is listed in the Apps/App Catalog.

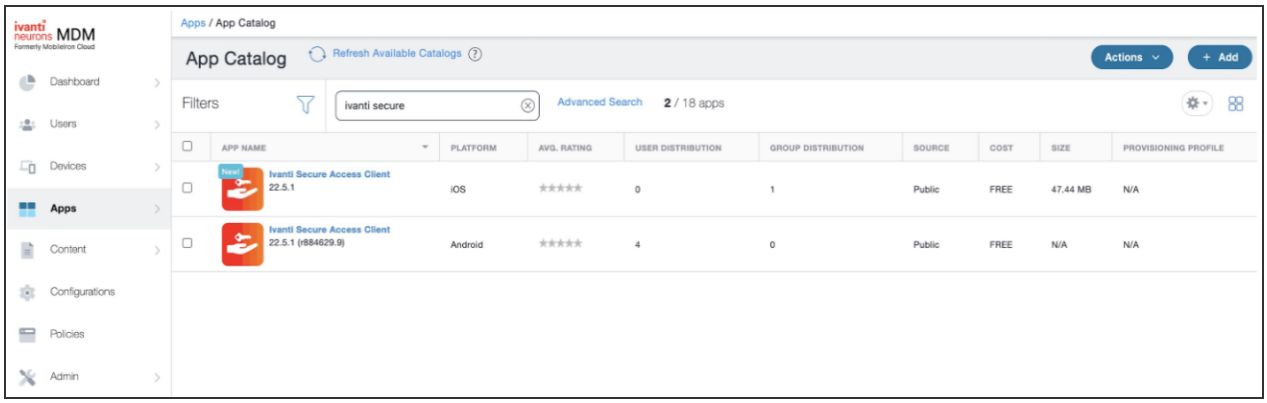

Integrating Ivanti NMDM with nZTA

To link Ivanti NMDM with nZTA:

1. From the nZTA menu, select **Integrations** > **MDM**.

   The MDM Configuration page appears.

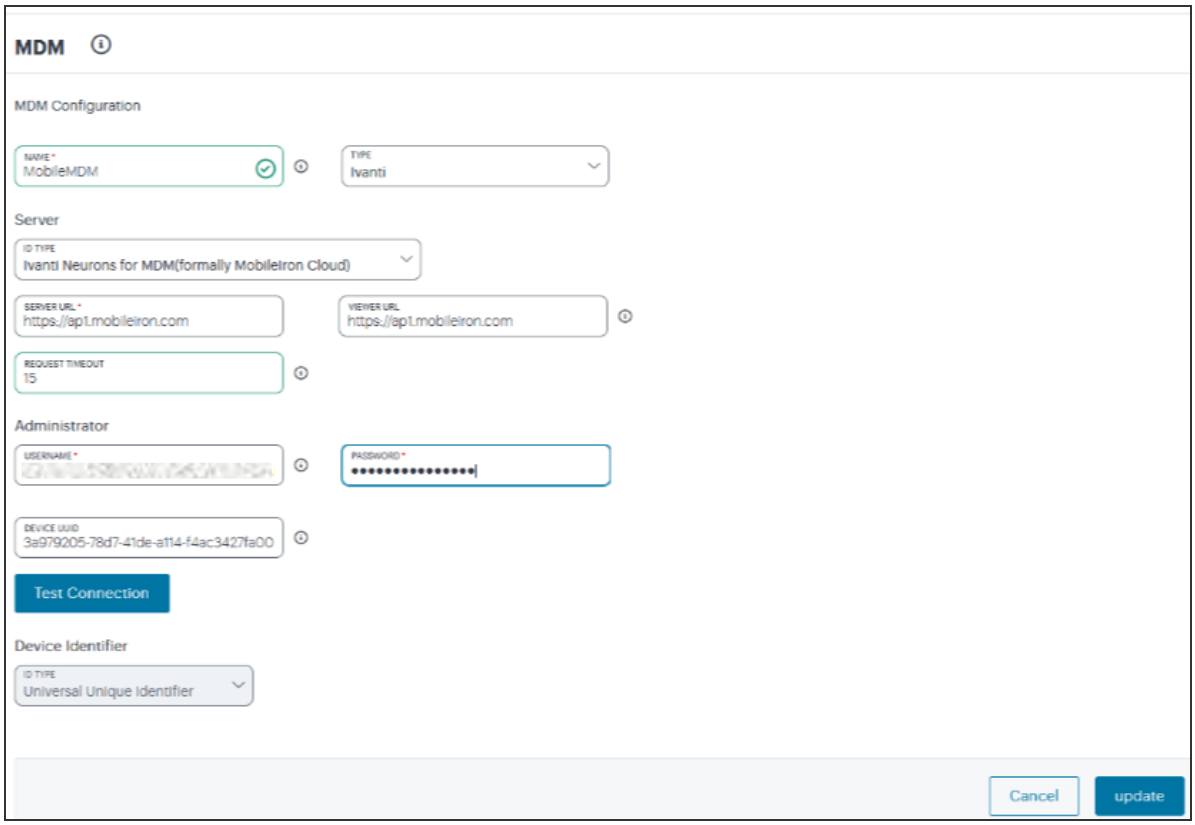

2. Specify a unique name for the MDM configuration.
3. Select **Type** as *Ivanti* (default).
4. Select **Server Type** as *Ivanti Neurons for MDM* (default), and enter its URL for connecting to the server.
5. Specify the URL for the MDM report viewer.
6. Specify a timeout period (0-60 seconds) for queries to the MDM server.
7. Specify the Administrator user name and corresponding password that has privileges to access MDM RESTful Web API.

8. Enter your **Device UUID**, and click **Test Connection**. A pop up shows the supported device attributes.

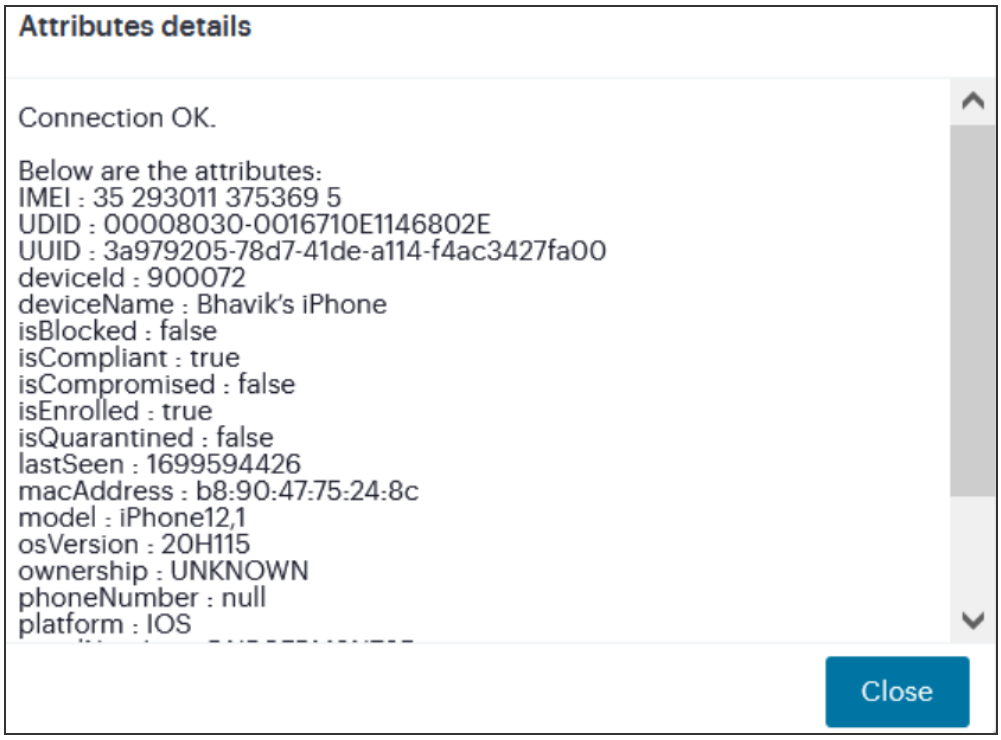

9. The **Device Identifier Type** for MDM certificate configuration is set as *Universal Unique Identifier*.

10. Click **Create**.

Creating Device Policy

You can create Device policies and then create / associate one or more Device Rules as required.

To configure device policy:

1. Select **Secure Access**, and then select **Manage Devices > Device Policies**.

2. Click **Create Device Policy**.

3. Enter policy name and description.

4. Click **Create Device Rule**.

   The Create Device Rule dialog appears.

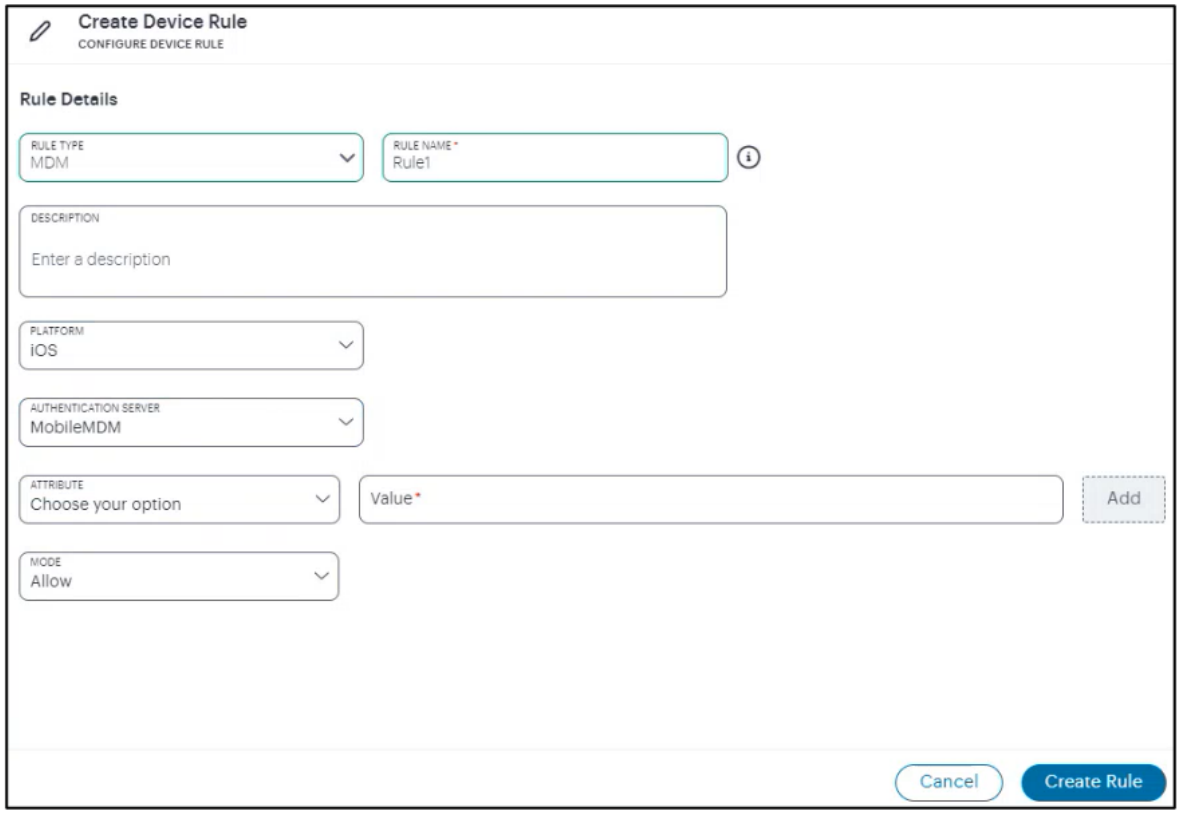

5. Select **MDM** as Rule Type.
6. Enter rule name and description.
7. Select **iOS** from the **Platform** drop-down list.

   The 22.6R1.2 release supports only iOS platform.

8. Select **MobileMDM** as Authentication Server.

9. Select a device attribute from the drop-down list, then enter a matching value for the selected attribute, and then click **Add**.

   Repeat the step to add more device attributes. The list of supported device attributes:

deviceId	deviceName	isBlocked	isCompliant	isCompromised
isEnrolled	isQuarantined	IMEI	lastSeen	macAddress
model	osVersion	ownership	phoneNumber	platform
serialNumber	UDID	UUID	userId	userEmail
userName				

   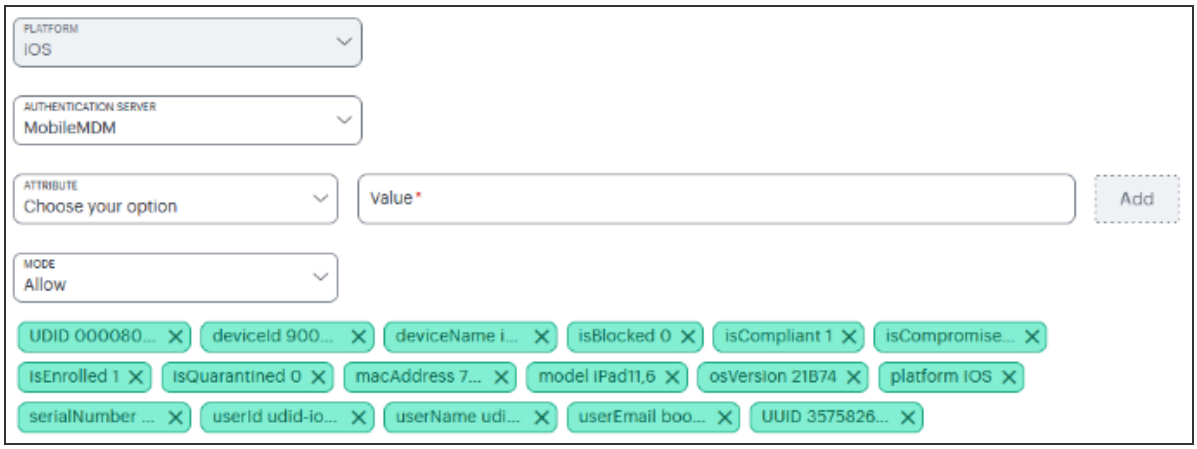

   - The lastSeen attribue is supported with the unix timestamp.
   - The isBlocked/isEnrolled/isCompliant/isQuarantined/isCompromised attributes support only 0(false) and 1(true) values in the configuration.

   To learn more about creating device rules, see "Creating Device Policy Rules" on page 451.

10. Select **Mode** and select one of the following from the drop-down list.

    - **Allow**: Enforces the policy rule.
    - **Deny**: Denies the policy rule.

11. Click **Create Rule**.

12. In the Device Policies page, click **Create Device Policy**.

Creating User Policy

To create a user policy:

1. From the *nZTA* menu, select **Secure Access** and then select **Manage Users > User Policies**.
2. Click **Create User Policy**. For details, see "Configuring Default Device Policy for Users" on page 448.

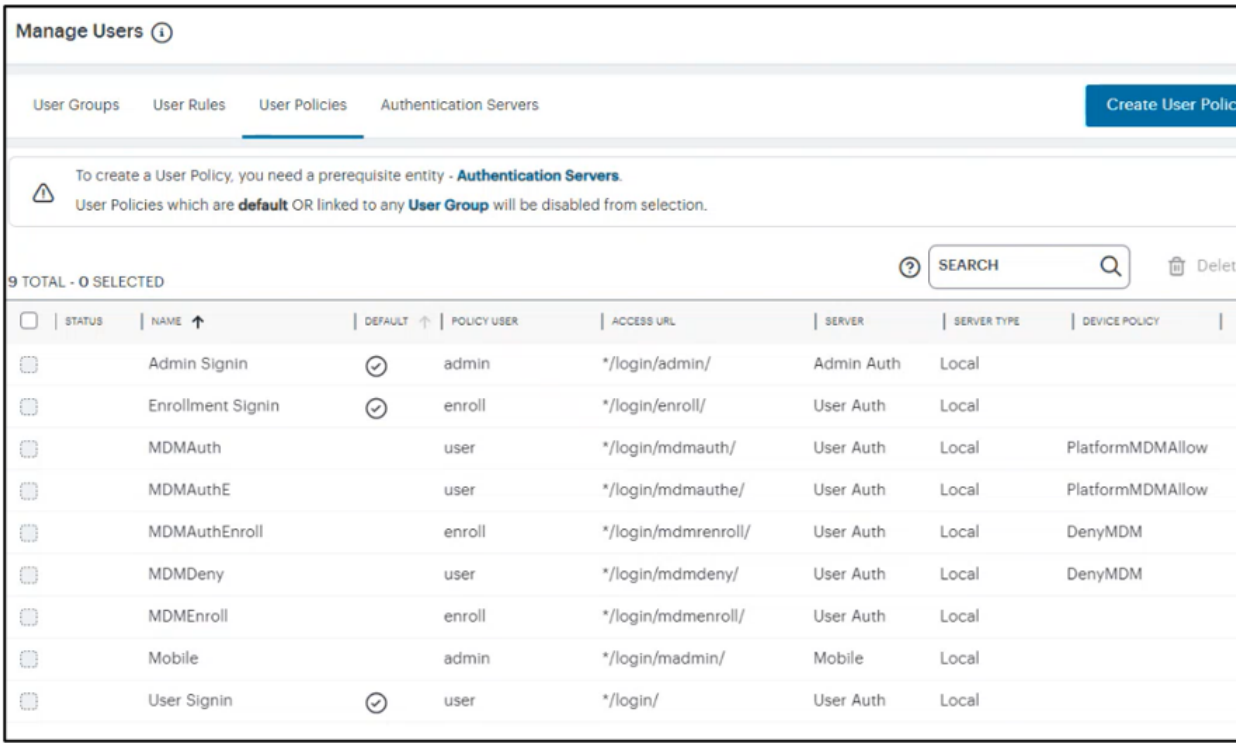

Manage User Policy

Creating Secure Access Policy

To create a secure access policy:

1. From the *nZTA* menu, select **Secure Access > Secure Access Policies**.

   The **Secure Access Policies** page appears. This lists all current secure access policies.

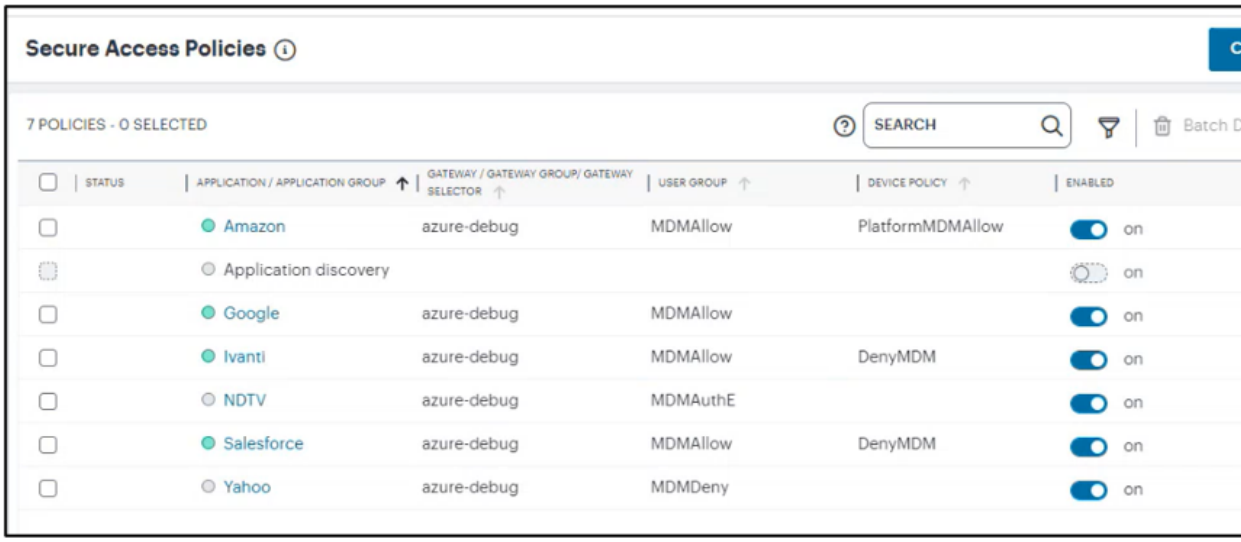

   Creating Secure Access Policies page

2. Click **Create**. For details, see "Creating a Secure Access Policy" on page 504.

End User Experience

Ivanti Neurons for MDM provides compliance check and simplified onboarding experience for nZTA end users connecting via mobile.

- When MDM policies are not enforced, both Enrollment and Authentication are successful.

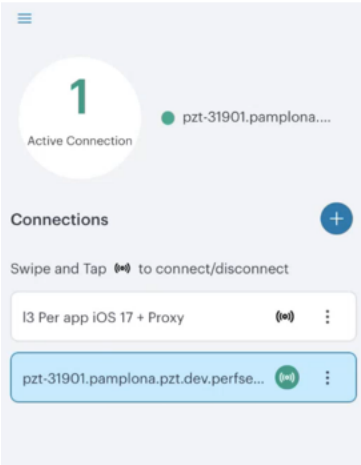

- If MDM policy is Deny for Enrollment, then Enrollment fails.

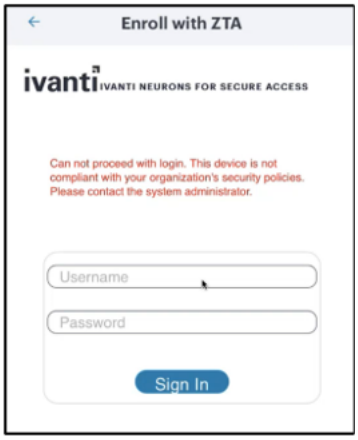

- If MDM policy is Allow for Enrollment, but Deny for Authentication, then Enrollment is successful but Authentication fails.

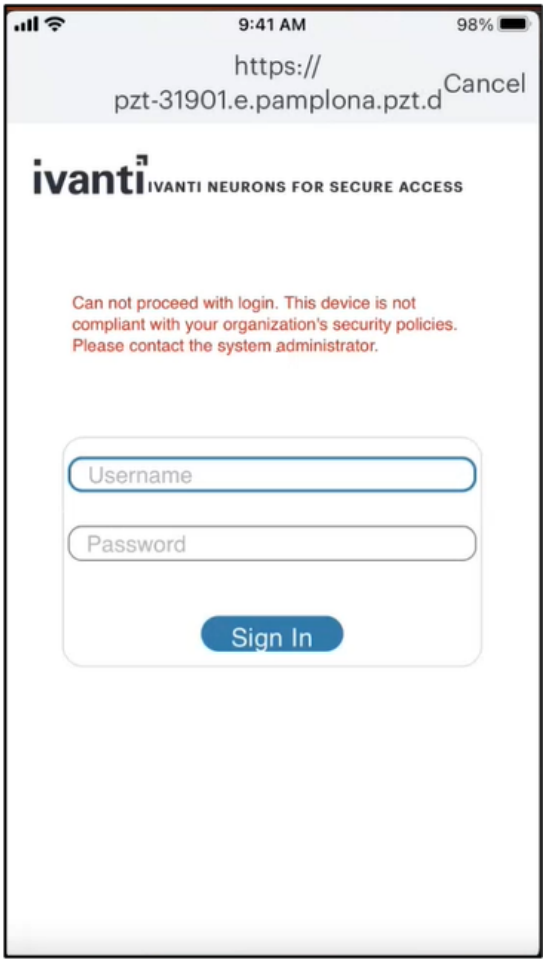

- If MDM policy for a certain application is Deny application access, then access to that application is blocked.

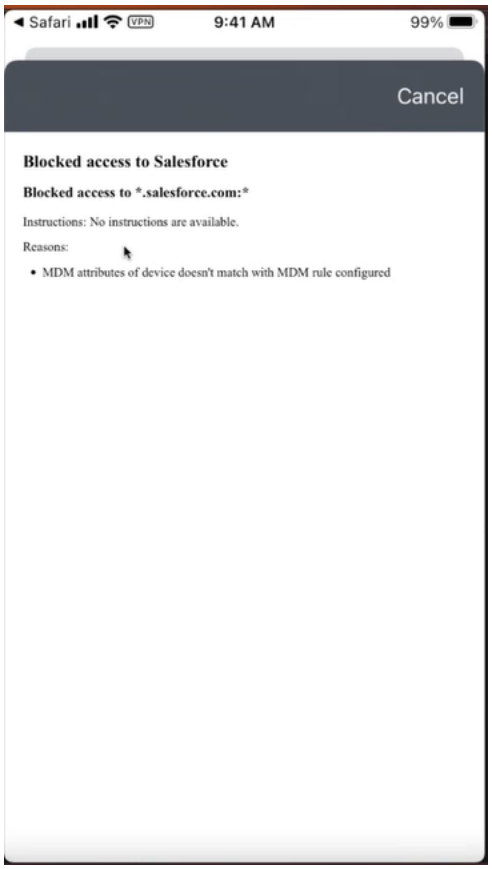

Working with User Authentication

Introduction

After you have logged into the *Controller* for the first time (see Logging in as a Tenant Administrator), you can create **authentication methods** and apply them to the **authentication policies** you define in your *Ivanti Neurons for Zero Trust Access* (*nZTA*) deployment. You then apply an authentication policy, together with **user rules**, to a **user group**. A user group forms part of a *Secure Access Policy*.

> To learn more about secure access policies, see Creating/Editing Secure Access Policies.

To view user authentication methods currently defined on the *Controller*, see Viewing User Authentication Methods. To view user authentication policies, see Viewing User Authentication Policies.

This chapter includes workflows for configuring user authentication according to each supported authentication type. *nZTA* supports the following types:

- **Local authentication**: An authentication system that is internal to the *Controller*. You must create all users manually on the *Controller*, and update any required authentication policies. see Workflow: Creating a Local Authentication Policy.

- **Azure AD SAML authentication**: An existing remote SAML authentication system based on an Azure AD server. See Workflow: Creating a SAML Authentication Policy With Azure AD.
- **On-premICSSAML authentication**: An existing remote SAML authentication system based on an on-premises *ICS* server. See Workflow: Creating an Authentication Policy for On-Premises ICS SAML.
- **Okta SAML authentication**: An existing remote SAML authentication system based on Okta. See Workflow: Creating a SAML Authentication Policy for Okta.
- **PingID SAML authentication**: An existing remote SAML authentication system based on PingID. See Workflow: Creating a SAML Authentication Policy for Ping Identity.
- **TOTP authentication**: Time-based One Time Password (TOTP) authentication as a secondary mechanism in *Multi-Factor Authentication* deployments. See Workflow: Adding TOTP to an Authentication Policy.

After you have created the required authentication methods and updated your user authentication policies, you create user rules and user groups, see Creating User Rules and User Groups.

Optionally, you can associate each user group with an *admin role*, see Associating User Groups with Admin Roles.

Viewing User Authentication Methods

To view the user authentication methods defined on the *Controller*, select the **Secure Access** icon in the *nZTA* menu, then select **Manage Users > Authentication Servers**.

The *Authentication Servers* page appears, showing all user authentication methods:

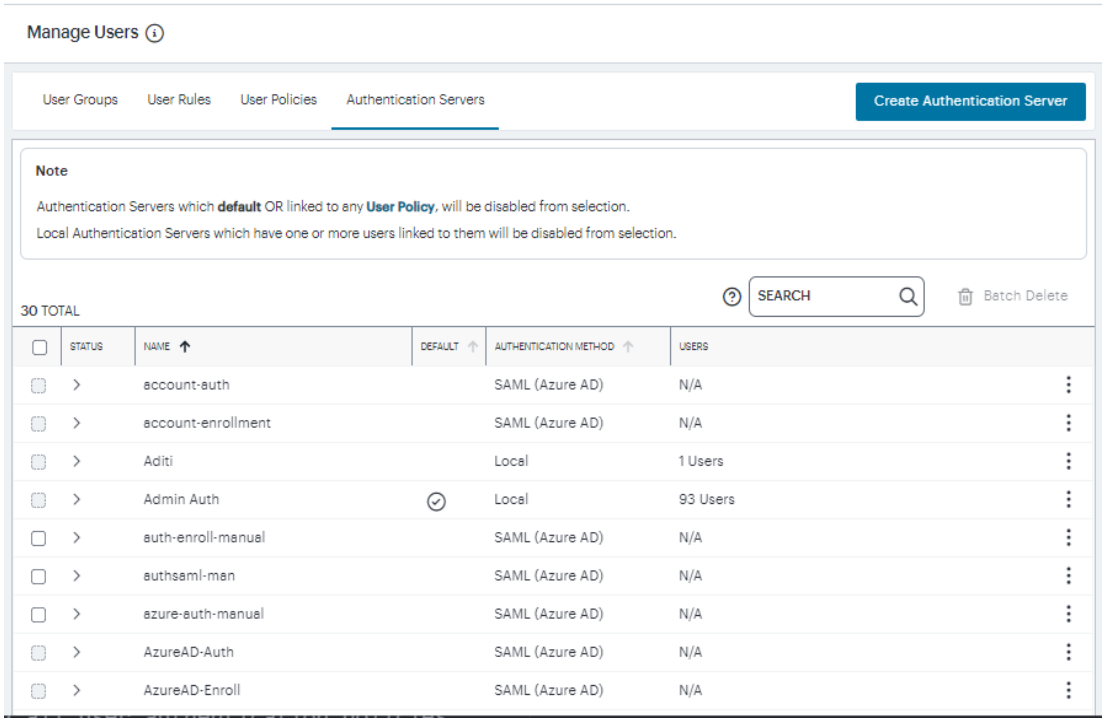

User Authentication Methods

From this page, you can:

- Add a new authentication method by selecting **Create Authentication Server**.
- Edit an existing authentication method by clicking the adjacent three dots, then selecting **Edit**. Make any required updates and save the changes.
- Delete an unused authentication method by clicking the adjacent three dots, then selecting **Delete**. You must confirm the deletion.
- View the configured attributes for a SAML authentication method, where that method is configured for use with an authentication policy. To do this, select the arrow indicator to the left of the method name, where shown.

Viewing User Authentication Policies

To view the user authentication policies defined on the *Controller*, select the **Secure Access** icon in the *nZTA* menu, then select **Manage Users > User Policies**.

The *User Policies* page appears, showing all user authentication policies.

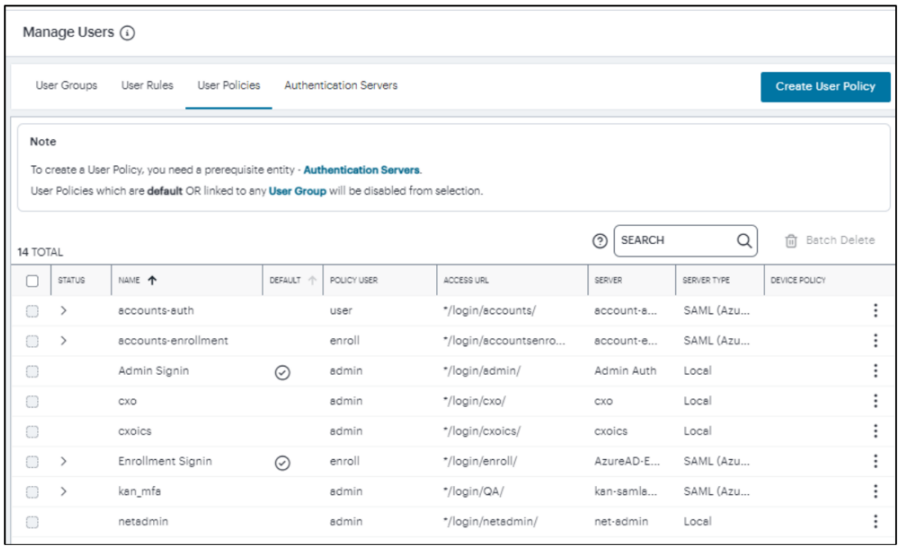

User Authentication Policies

nZTA provides following default/built-in authentication policies, indicated by a tick in the **Default** column, suitable for the primary use-cases of administrative sign-in, user enrollment, and user sign-in:

- *Admin SignIn*. This policy is used whenever admin users log in. That is, for connection requests to the **/login/admin/* URL. It is referenced by the *ALLADMINUSERS* user rule, which associates it with the *ADMINISTRATORS* user rule group.
- *User SignIn*. This policy can be used as the primary connection endpoint for all user device sign-in and enrollment requests. That is, for connection requests to the **/login/* URL. It is referenced by the *ALLUSERS* user rule, which associates it with the *USERS* user rule group.a

These policies are fixed and cannot be deleted. However, you can edit them to reference specific authentication methods.

Furthermore, you can create additional custom authentication policies to enable bespoke authentication for specific groups of users or parts of your organization. Each policy should contain a unique access URL to which your users connect, and each should then be configured to link to authentication methods applicable for that purpose (for more information, see Adding Custom Authentication Policies).

To learn more about how user authentication policies are used in a *nZTA* service, see Defining User Authentication.

From this page, you can:

- (For SAML authentication) Download policy metadata files that are required for external SAML enrollment or sign-in apps. To do this, select the check box for the required policy and select **Download**. Save the file to your local workstation.
- View the configured attributes for a SAML-authenticated policy, where that policy is configured with a valid SAML authentication method. To do this, select the arrow indicator to the left of the policy name, where shown.
- Add an authentication policy by selecting **Create User Policy**.
- Edit an existing authentication policy by clicking the adjacent three dots, then selecting **Edit**. Make any required updates and save the changes.
- Delete an unused authentication policy by clicking the adjacent three dots, then selecting **Delete**. You must confirm the deletion.

Creating User Rules and User Groups

After your authentication method is established and associated with an authentication policy, you can set up any required *user rules* and *user groups*. A user rule identifies one or more users based on a test against a selected attribute present in a user credential or profile, checked against either a local authentication record or from a SAML authentication service. For information about creating user rules, see Creating User Rules.

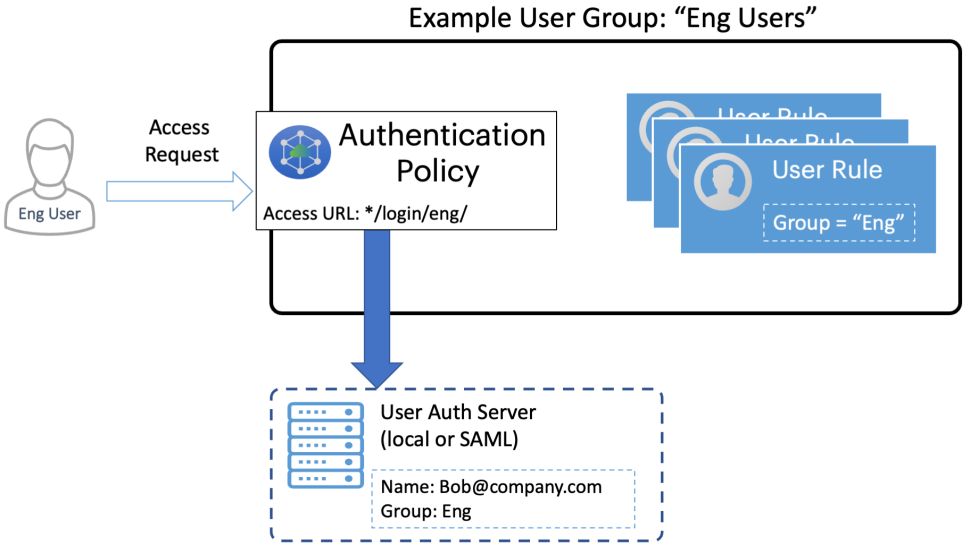

Performing authentication through a User Group

You associate one or more user rules with an *authentication policy* to form a *user group* (see Creating User Groups). Users requesting authorization for a service controlled by a Secure Access Policy must pass all the rules contained in the User Group attached to the policy.

A user group is required when defining a *secure access policy*. The user group identifies the users and the authentication policy to which a secure access policy applies, see Creating/Editing Secure Access Policies.

Optionally, you can associate each user group with an admin role, see Associating User Groups with Admin Roles.

Creating User Rules

Through user rules, an admin can construct a test to provide authorization to only those users of a particular name, role, group, or some other stored attribute. In the rule configuration, you select the user attribute on which you want a test to be performed.

nZTA includes following default user rules:

- *ALLADMINUSERS*. This matches all users, and is referenced by the default *ADMINISTRATORS* user group, which associates it with the built-in *Admin Signin* authentication policy.
- *ALLUSERS*. This matches all users, and is referenced by the default *USERS* user group, which associates it with the built-in *User Signin* authentication policy.

> To read more about default user groups, see Creating User Groups. To read more about built-in authentication policies, see Viewing User Authentication Policies.

This preset configuration of rules, groups, and policies is suitable for typical use cases involving whole-organization authorization needs. In other words, where you require only a single user authorization path that matches all users. For scenarios where you require more specific user authorization checks, you can create additional rules to match specific types of users.

When you create a rule, you select the user attribute with which you want this rule to test. *nZTA* provides the following rule attribute types:

- **username**: For local authentication methods, choose this attribute type to match against locally-defined user names.
- **SAML (Azure AD)**: For SAML authentication methods, choose this attribute type to match against user names or groups provided by the SAML service.

- **Custom**: For SAML authentication methods, choose this attribute type to match against a custom SAML attribute expression.

To create a user rule:

1. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > User Rules**.

   The *User Rules* page appears. This page lists all user rules.

2. Click **Create User Rule**.

   The *Create User Rule* form appears.

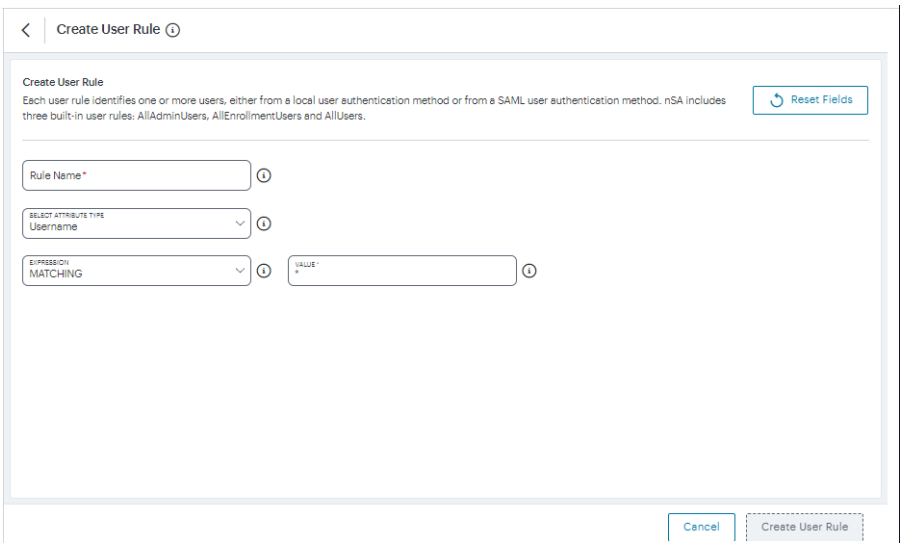

Create User Rule

At any point during this process, you can reset the form data by clicking **Reset Fields**.

3. Enter a **Rule Name**.

4. Select **Select Attribute Type** and select one of the available options:

    - *Username*: Matches user names in a local authentication method. When you select this option, you must then:

        - Select an **Expression** type, either *Matching* or *Not Matching*.
        - For the **User** value, enter a match expression for the selected **Expression** type. For the value:

            - A comma-separated list of items is supported where required.
            - Wildcard matches are supported.
            - Special characters are supported.
            - Single and double quotes are not supported.

        Ivanti recommends that a basic asterisk wildcard is not used when you intend to associate admin roles with user groups. Instead, a more-specific wildcard that only includes admin users is required in this case to prevent all users having total access rights.

    - *SAML (Azure AD)*: Matches user names or groups in a SAML authentication method. When you select this option, you must then:

        - Select a **SAML Attribute Type**, either *Username* or *Group*.
        - For **Attribute Value**, enter a single-value match expression for the selected **SAML Attribute Type** as a SAML expression. For matches against *username*, *-wildcard values are accepted. For matches against *Group*, specify either a Group Name or GroupID based on the configuration in your Azure AD service.

        If you select an attribute type of *Group*, an authenticated user must be a part of only the matching group. If the user is a part of multiple groups, authentication will fail.

- *Custom*. Matches against a custom SAML attribute expression. When you select this option, use the **Type or Create an Expression** property to enter an attribute expression. Supported formats include:

  - For simple user attribute key-value matching, use the syntax `userAttr.<attr-key> [=|!=] <attr-value>`. For example:

    ```
    - userAttr.memberOf = "CN=sales,DC=example,DC=com"
    - userAttr.mail = "user1@example.com"
    - userAttr.realm = "Users"
    - userAttr.department != "example_department"
    ```

  - To match against attributes that can have multiple values associated with a single attribute key, use the syntax `samlMultiValAttr.<attr-key> [=|!=] (<list>)`. For example:

    ```
     - samlMultiValAttr.memberOf =
    ("CN=Employee,CN=Users,DC=example_demo,DC=com")
    - samlMultiValAttr.memberOf = ("CN=Users,DC=example_
    demo,DC=com")
    ```

  - Use brackets and AND/OR operators to construct logical compound expressions:

    ```
     - userAttr.groups = ("Group1" or "Group2")
    - userAttr.realm = ("ztaqa") and samlMultiValAttr.memberOf =
    ("CN=sales,DC=uisdp,DC=com")
    - userAttr.realm = ("ztaqa") or samlMultiValAttr.memberOf =
    ("CN=sales,DC=uisdp,DC=com")
    - userAttr.realm != ("ztaqa") and samlMultiValAttr.memberOf
    = ("CN=sales,DC=uisdp,DC=com")
    ```

  

  For *samlMultiValAttr* expressions containing multiple groups (for example, `samlMultiValAttr.groups = ("Group1" or "Group2")`), your matching users must be a part of both groups in the expression to obtain authorization.

5. To create the user rule with the displayed settings, click **Create User Rule**.

   The new user rule is added to the list of user rules.

6. Repeat steps 3-6 for each required user rule.

7. (Optional) Edit an existing user rule by clicking the adjacent three dots, and then selecting **Edit**. Make any required updates and save the changes.
8. (Optional) Delete an unused user rule by clicking the adjacent three dots, and then selecting **Delete**. You must confirm the deletion.

After you have created all required user rules, you can create user groups, see Creating User Groups.

Creating User Groups

After you have created user rules (see Creating User Rules), you associate one or more user rules with an authentication policy to form a user group.

> User groups are one of the four dimensions of a Secure Access Policy, see Creating/Editing Secure Access Policies.

nZTA includes following default user groups:

- *ADMINISTRATORS*. This user group associates the default *ALLADMINUSERS* user rule with the built-in *Admin Signin* authentication policy.
- *USERS*. This user group associates the default *ALLUSERS* user rule with the built-in *User Signin* authentication policy.

> To read more about built-in authentication policies, see Viewing User Authentication Policies.

This preset configuration of rules, groups, and policies is suitable for typical use cases involving whole-organization authorization needs. In other words, where you require only a single user authorization path that matches all users. For scenarios where you require more specific user authorization checks, you can create additional user groups to make different associations of user rules and custom authentication policies.

To create a user group:

1. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > User Groups**.

   The *User Groups* page appears. This page lists all user groups.

2. Click **Create User Group**.

   A form appears to enable you to create the user group.

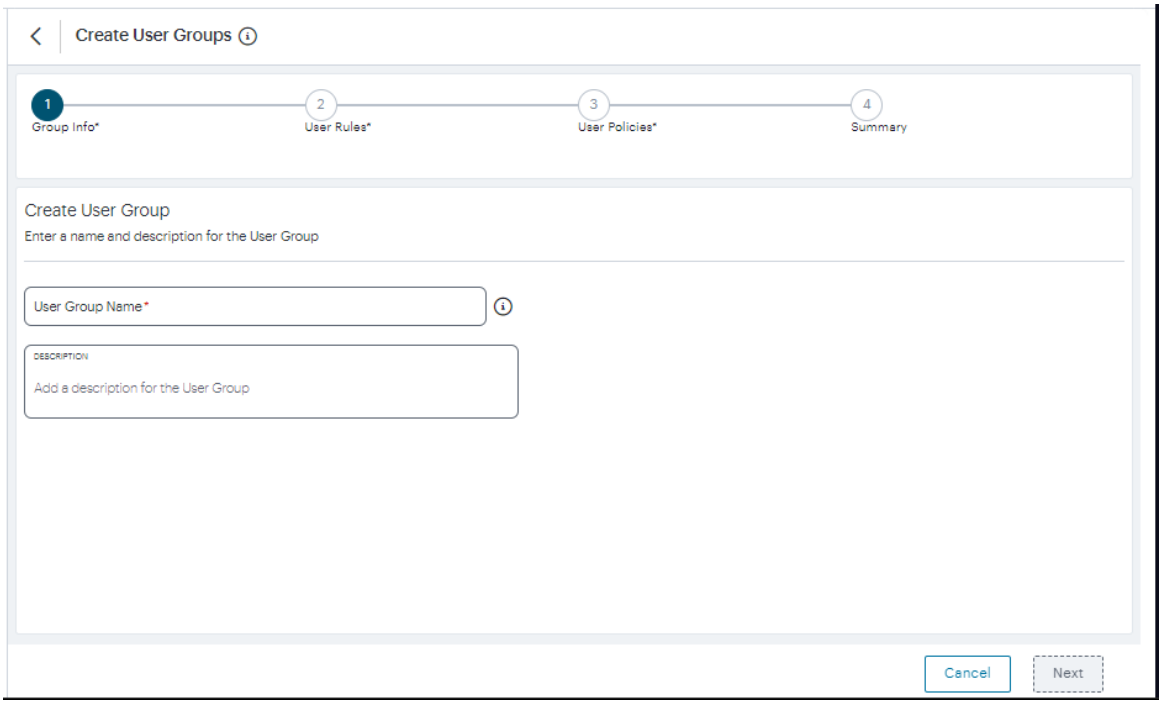

Create User Groups

3. Enter a **User Group Name** and an optional **Description**, then click **Next**.
4. Select each of the listed **User Rules** that are required in the user group, then click **Next**.

   User rules should be different for enrollment and user sign-in groups

5. Select required authentication policy from the list, then click **Next**.
6. Review the summary and click **Create**.

   The new user group appears in the **User Groups** list.

7. Repeat steps 2-7 to create all required user groups.
8. (Optional) To edit a listed user group, click the adjacent three dots, then select **Edit** and make any required updates.
9. (Optional) To delete an *unused* user group, click the adjacent three dots, then select **Delete** and confirm the deletion.

After you have created user groups, you can optionally assign the user group to an admin role, see Associating User Groups with Admin Roles.

Associating User Groups with Admin Roles

An admin role defines the elements of the user interface that an associated user group can access.

The current user can only access an individual user interface page/workflow if their user group is associated with an admin role that permits it. The tasks they can perform within that displayed element depends on the permissions set within the admin role.

When you are using admin roles, *Ivanti* recommends that any user rules for administrators does not use a basic asterisk wildcard, see Creating User Rules. Instead, a more-specific wildcard that only includes admin users is required in this case to prevent all users having total access rights.

The default admin roles are not created by the tenant admin using the *nZTA* user interface. Rather, they are set up by the *Ivanti* DevOps team.

For example, the DevOps team might define the following admin roles:

- The *.Administrators* admin role has access to all user interface elements (full read, create, update, delete rights).
- The *.Read-Only Administrators* admin role has access to all user interface elements except workflows (read only).
- The *.Network Administrators* admin role has access to *nZTA Gateways* and Insights (read only).
- The *.CxOs* admin role has access to Insights only (read only).

  For more information about your assigned admin roles, please contact *Ivanti* DevOps.

The Tenant Admin can view admin roles in the **Administration > Admin Roles** page, and associate each role with a single user group.

To associate a user group with an admin role:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, select **Administration**, then select **Admin Roles**.

   A list of Admin Roles appears. This includes default admin roles and custom admin roles (RBAC). For example:

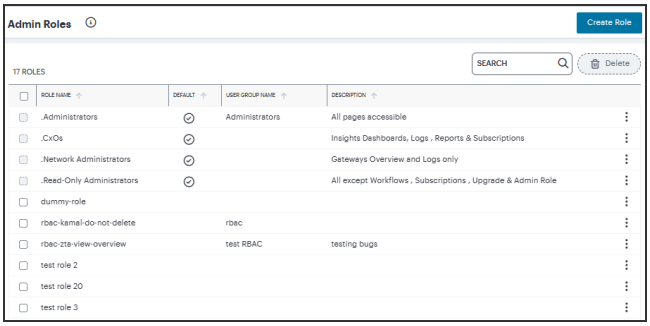

   Admin Roles

3. Click the three dots adjacent to the role you want to update, then select **Edit Role**.

   A dialog appears. For example:

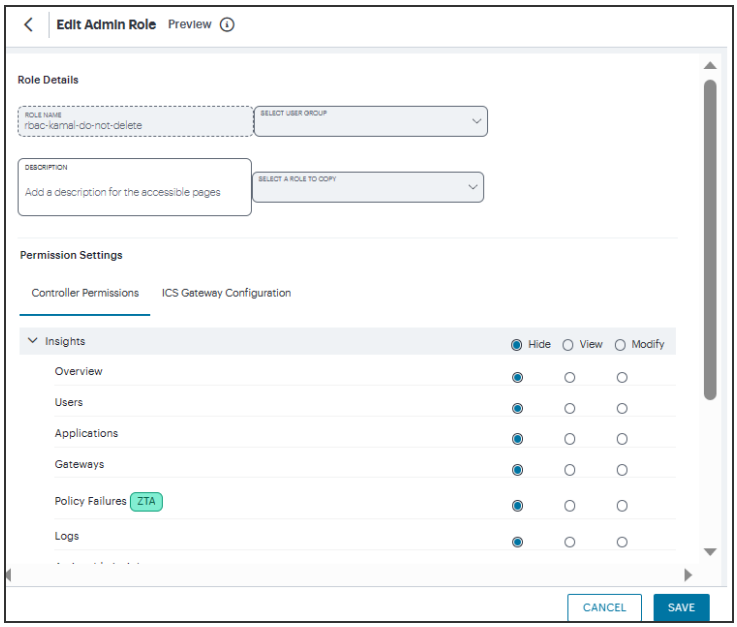

   Edit Admin Roles

4. Under **Choose group**, select the user group that you want the admin group to be associated with.
5. Select **Save Changes**.

6. (Optional) Repeat steps 3 to 5 for each admin role.

Role-based Access Control for Admin Users

With Role-based access control (RBAC), organizations can easily add admins and assign them specific roles, with differing levels of access to the nSA Admin Portal. In addition to an existing set of default roles, Administrators can now create custom granular roles for specific functions within the nSA admin portal.

The following examples illustrate how an organization can leverage role-based administration for a variety of scenarios.

To create a custom admin role:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, select **Administration**, then select **Admin Roles**.
3. In the Admin Roles page, click **Create Role**. The Create Admin Role page appears.

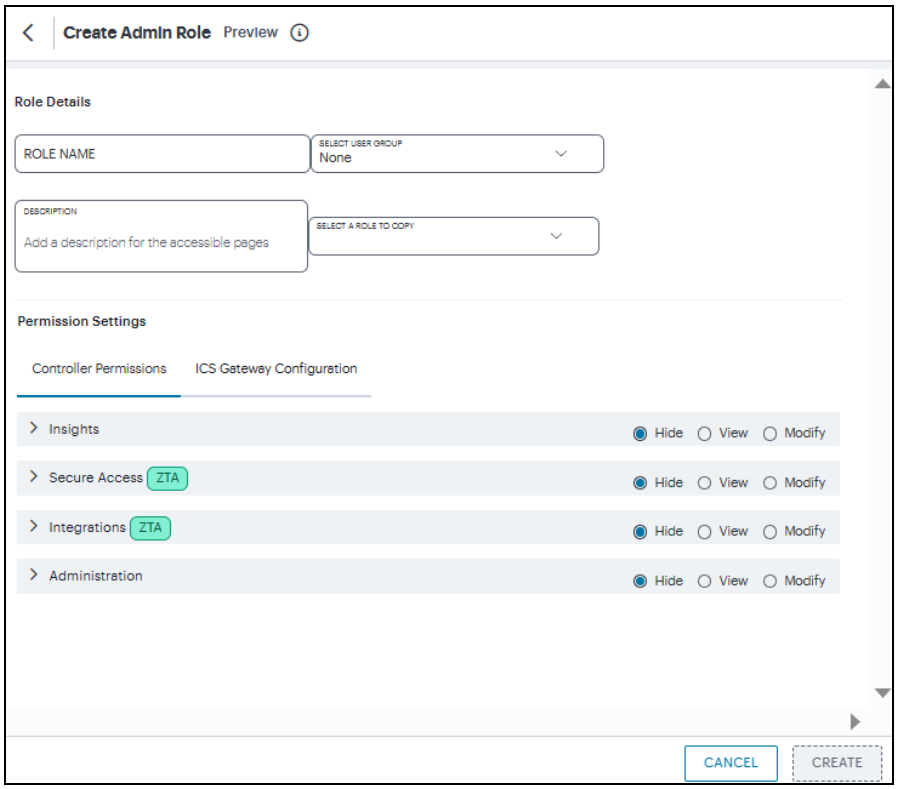

Create Admin Role

4. Enter a unique name for the role.

5. From the drop-down list, select the User Group that you want to associate with this role. For details, see "Associating User Groups with Admin Roles" on page 97.
6. Optionally, enter a **Description**.
7. From the drop-down list, select an existing role that suits your requirements.
8. Under Permission Settings, the *Controller* Permissions list shows the list of resources. The resources specific to nZTA are tagged with **nZTA** and resources specific to nSA are tagged with **ICS**.

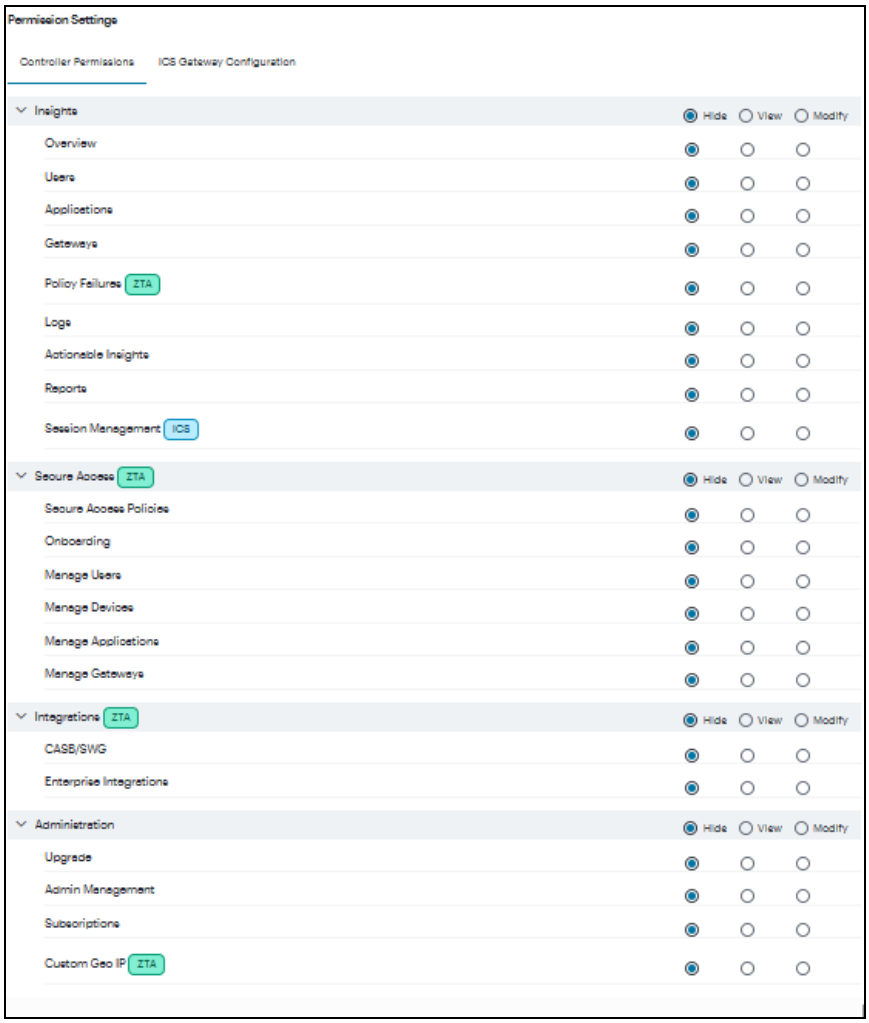

Controller Permissions

9. Select the Hide, View only, or Modify permissions for each resource and its attributes. This determines which pages to show and which actions to allow.
10. Under Permission Settings, the ICS Gateway Configuration list shows the list of ICS Gateway resources.

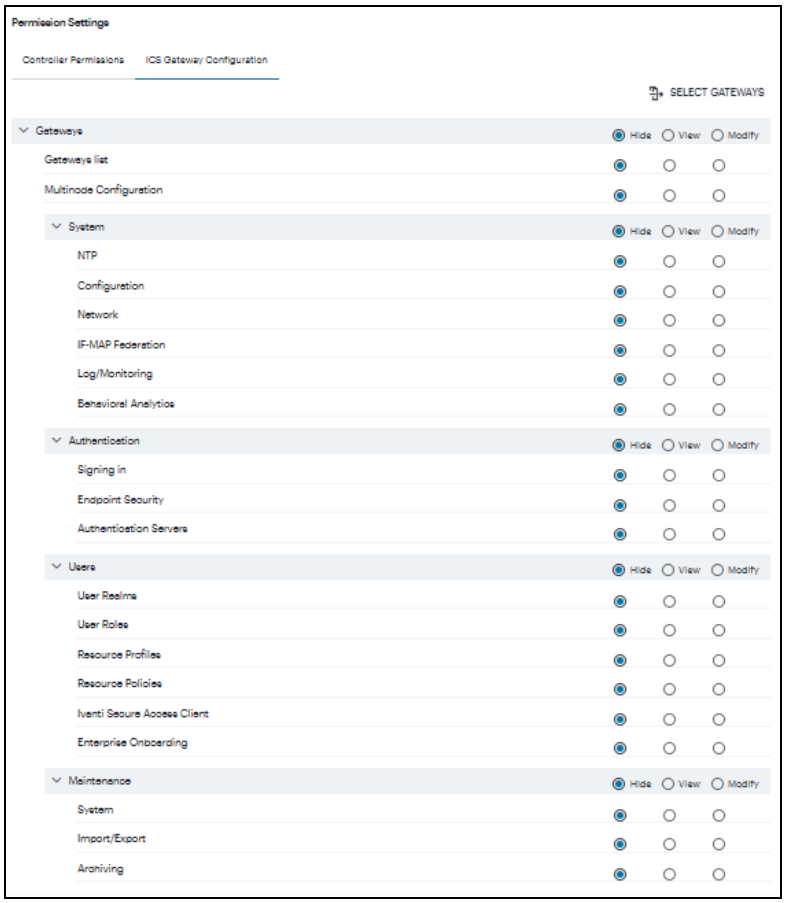

ICS Gateway Permissions

11. Click **Select Gateways**. In the Select Gateways dialog, select one or more Gateways / Clusters from the list and click **Apply**.

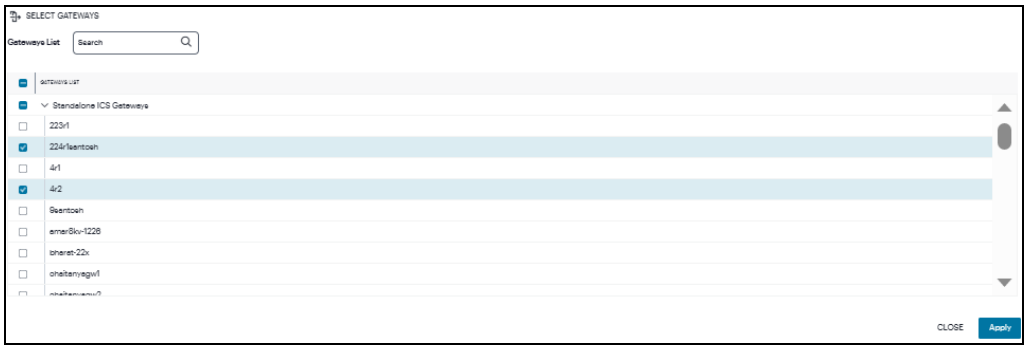

Select Gateways/Clusters

12. Select the Hide, View only, or Modify permissions for each resource and its attributes. This determines which pages to show and which actions to allow.

13. Click **Create**. The newly created custom admin role is displayed in the Admin Roles page.

14. (Optional) Edit an existing admin role by clicking the adjacent three dots, and then selecting **Edit**. Make any required updates and save the changes.

15. (Optional) Delete an unused custom admin rule by clicking the adjacent three dots, and then selecting **Delete**. You must confirm the deletion.

Workflow: Creating a Local Authentication Policy

This process involves creating a local user authentication *method* and defining within it all user credentials necessary to identify and authenticate your end-users. You then configure this method as the primary authentication method in your *authentication policies*. If you are configuring Multi-Factor Authentication (MFA) in your deployment, you can configure local user authentication as either the primary or secondary authentication method.

Before you begin, make sure you have all user details (name and password) ready.

To configure a *new* local authentication method:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.

2. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > Authentication Servers**.

   The **Authentication Servers** page appears. This page lists all existing user authentication methods. For example:

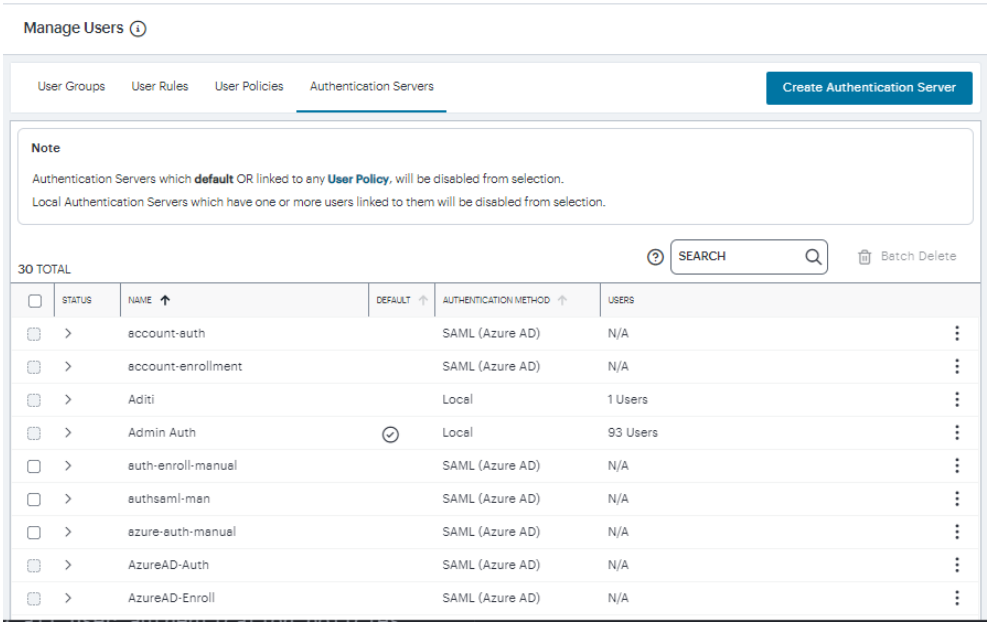

User Authentication Methods

3. Select **Create Authentication Server**.

   A form appears that enables you to define the authentication method.

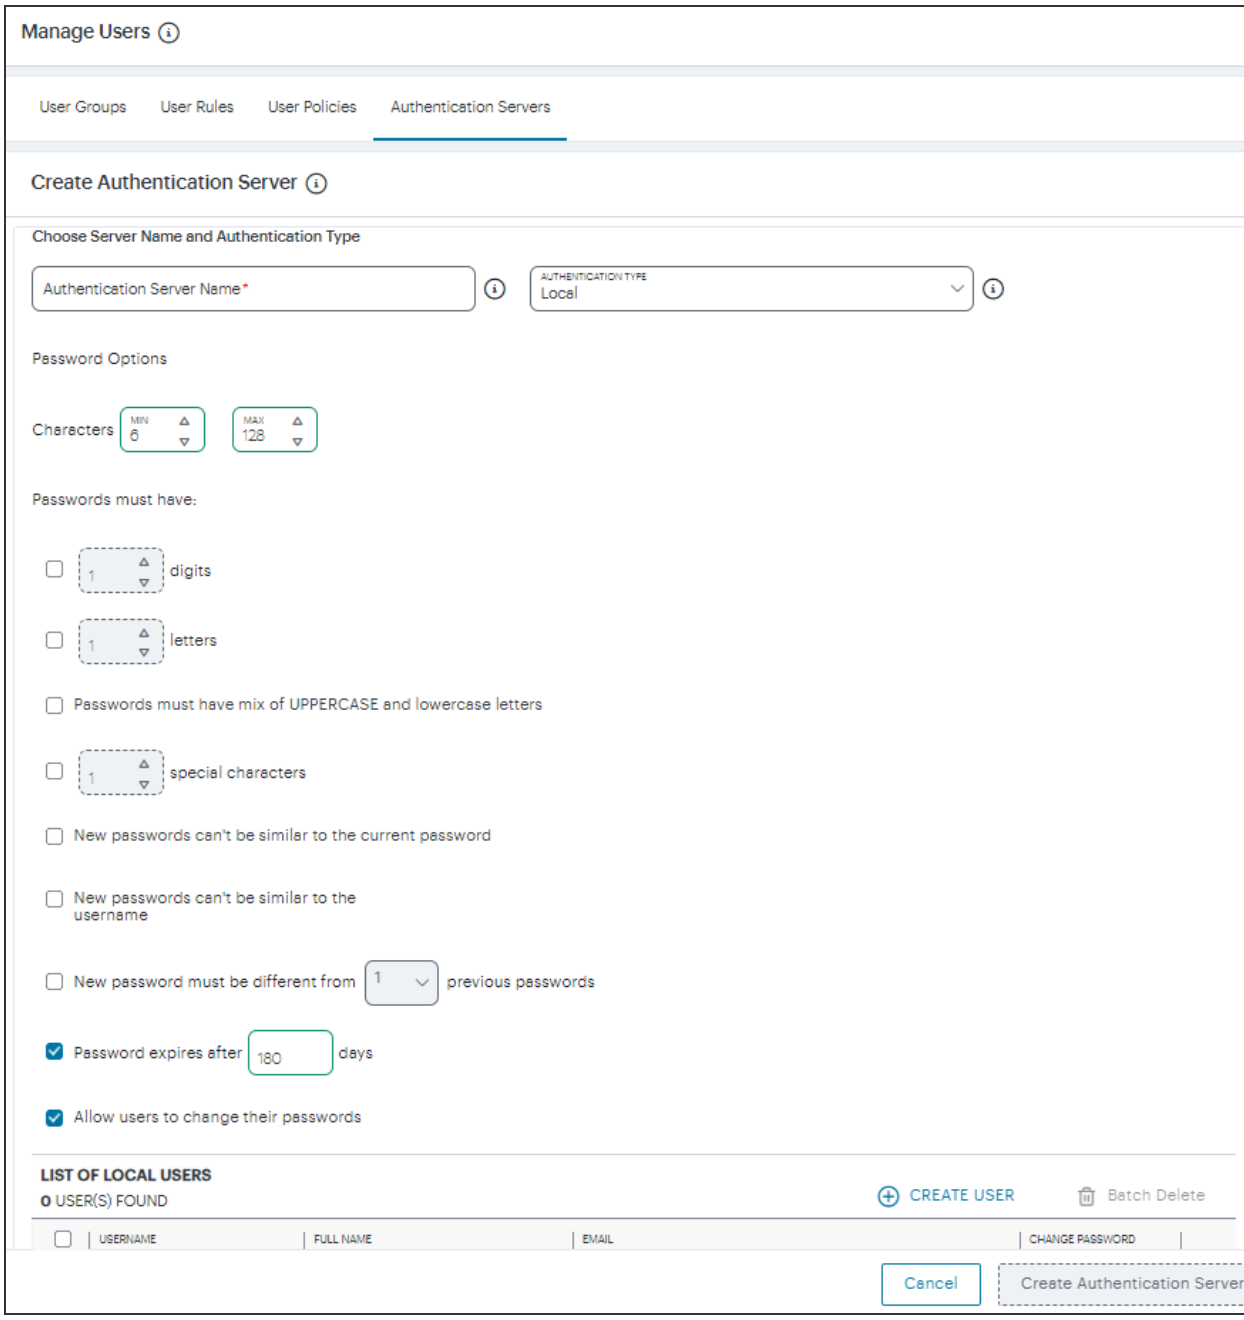

Adding a new local user authentication method

> At any point during this process, you can reset the form data by selecting **Reset Fields**.

4. Under **Choose name and type**:
    - Specify an **Authentication Server Name**.
    - Select the **Authorization Type** of *Local*.

5. Configure the password options. This is applicable to default Admin Auth and default User Auth.

Settings	Guidelines
Minimum length	Specify a number of characters. The valid range is 6-128. 6 is the default.
Maximum length	Specify a number of characters. The valid range is 6-128. 128 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.
Special Characters	Select this option if you want password should contain any special characters
Different from current password	Select this option if the password must not be same as the current password.
Different from username	Select this option if the password must not be same as the username.
Different from previous password	Select this option and then select the number from drop-down if a new password must not be same as the previous number of passwords.
Force password change	Select this option to specify the number of days after which a password expires. The default is 180 days.

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.
Prompt users to change password	Select this option to specify when to prompt the user to change passwords.

6. Click **Create User**. The *Create Local User* dialog appears to show additional local authentication settings:

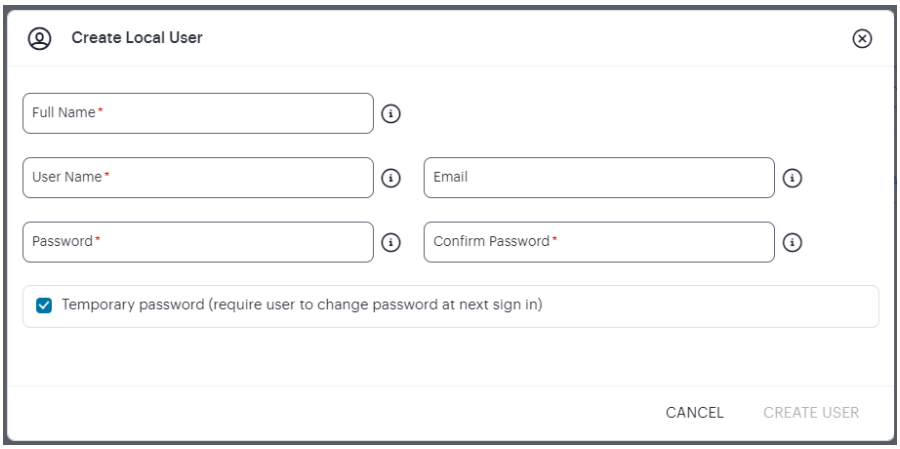

Adding local users to a new authentication method

7. Enter the following settings:

   - Specify a **User Name**, **Full Name**, and **Email** for the user.
   - Specify a **Password** and **Confirm Password** for the user.
   - (Optional) Select the **Temporary Password** check box if you want the user to change their password when they first log in.
   - Click **Create User**.

   The user is added to the list of users.

8. Repeat the previous step for each required user.

9. Click **Create Authentication Server**.

   The new local user authentication method is added to the list of methods and the process is complete.

After you have created your local authentication method, create or update your authentication policies with the new authentication method.

nZTA provides built-in policies to cover both basic cases. In addition, *nZTA* allows for the definition of custom policies to facilitate separate authentication endpoints for specific groups of users. To learn more, see Using User Authentication Policies.

Repeat the following steps for each policy, starting with enrollment:

1. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > User Policies**.

   The **User Policies** page appears. This page lists all existing user authentication policies.

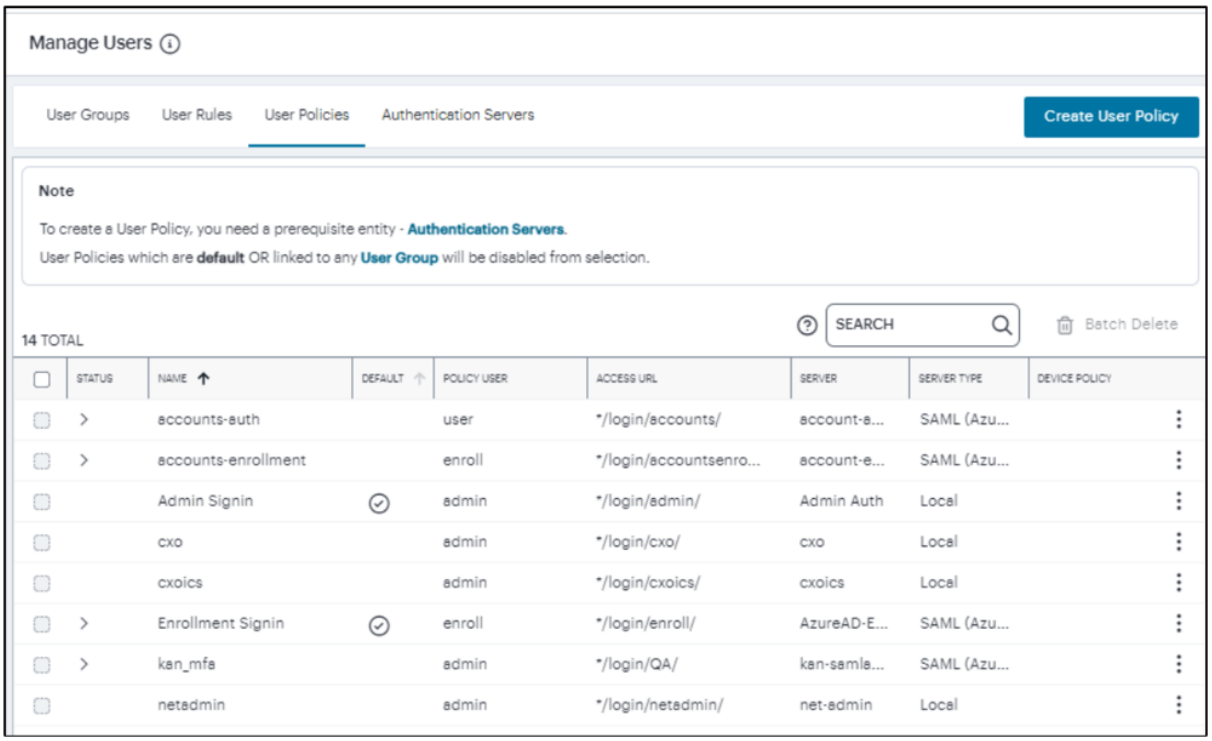

User Authentication Policies

   To learn more about the policies on this page, see Viewing User Authentication Policies.

   From this page, either create a new custom policy or edit an existing policy.

2. To add a new custom policy, click **Create User Policy**.

   The **Create Authentication Policy** form appears.

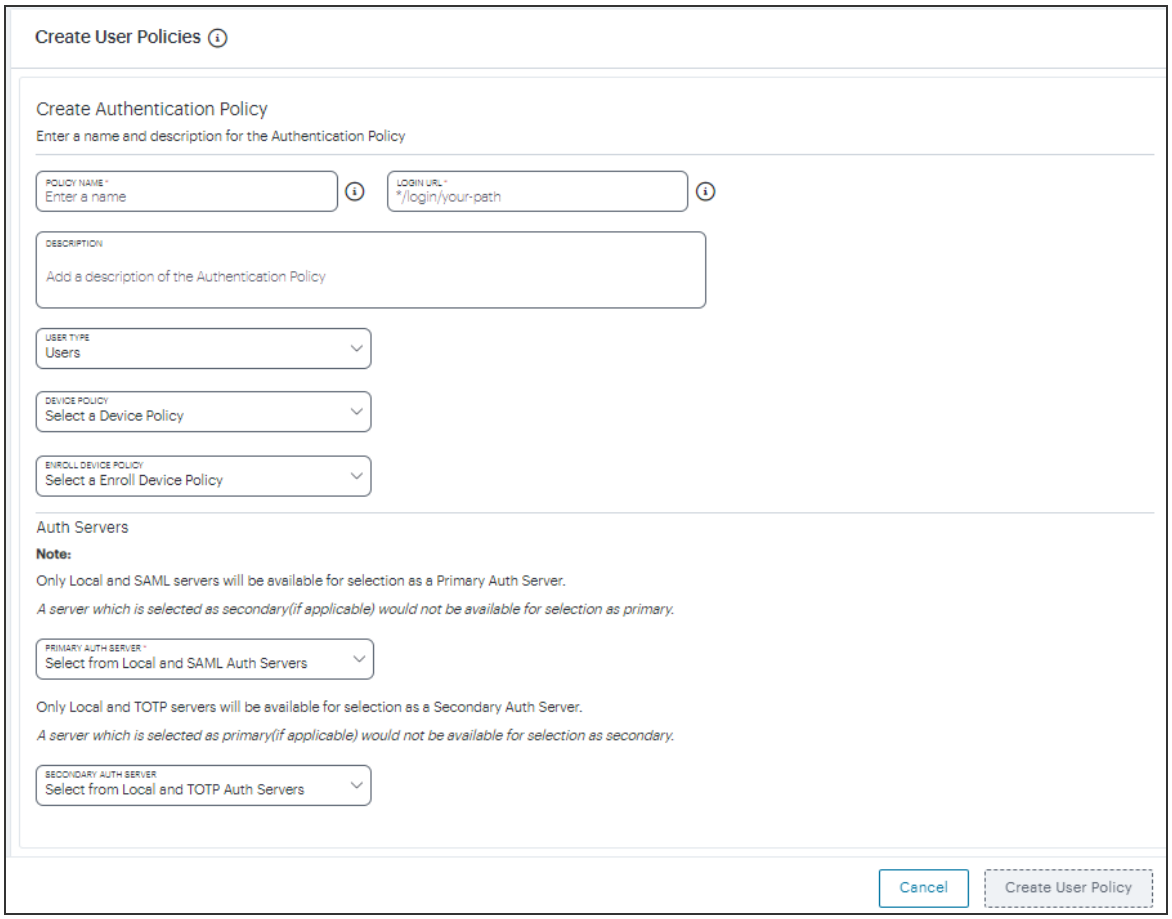

   Create Authentication Policy

   To learn more about how custom policies are used for user login and enrollment, see Adding Custom Authentication Policies.

3. Enter a **Policy Name**.

4. Enter a **Login URL** using the format **/login/<path>/*.

   The URL must start with "**/login/*" and cannot contain any special characters. <path> should be set to a unique value reflecting the endpoint URL you want to define for this authentication policy (appended with a backslash):

   - In the case of user sign-in policies, this is the URL endpoint (appended to the tenant FQDN) to which new users are invited to connect to enroll or sign-in a device with the *Controller*. Example value: "**/login/saleslogin/*".
   - In the case of user enrollment policies, this endpoint identifies the enrollment URL to which users are redirected if they attempt to connect to the equivalent sign-in policy with an un-enrolled device. In most cases, you do not advertise this endpoint to your users. Example value: "**/login/salesenroll/*".

   

   In some enrollment circumstances, such as when using a device pre-installed with an older version of *Ivanti Secure Access Client*, you connect directly to the enrollment policy endpoint to enroll the device. For more details, see Using User Authentication Policies.

5. (Optional) Enter a description for the authentication policy.
6. Select a **User Type** based on the intended authentication activity for this policy. Choose from:

   - **Users**: Select this option to define the user sign-in endpoint for enrolled devices. This is the endpoint that you provide to your users to access the service (regardless of enrollment status).
   - **Administrators**: Select this option to define the authentication endpoint for administrator-level sign-in. This endpoint is used for administrator login to the *Controller* only.

7. (for policies with a **User Type** of "Users" only): Select an **Enroll Device Policy** from the drop-down list to be linked to this sign-in policy (as indicated):

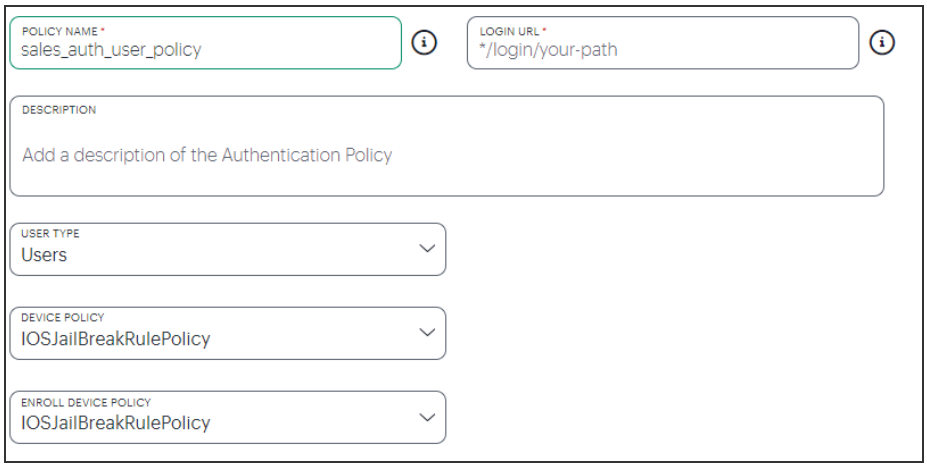

Linking an enrollment policy to a user sign-in policy

This is the enrollment policy to which a user is redirected if it is determined that the device is not yet enrolled. To learn more, see Using User Authentication Policies.

8. Under **Policy Server Details**, select **Primary Auth Server** from the drop-down list:

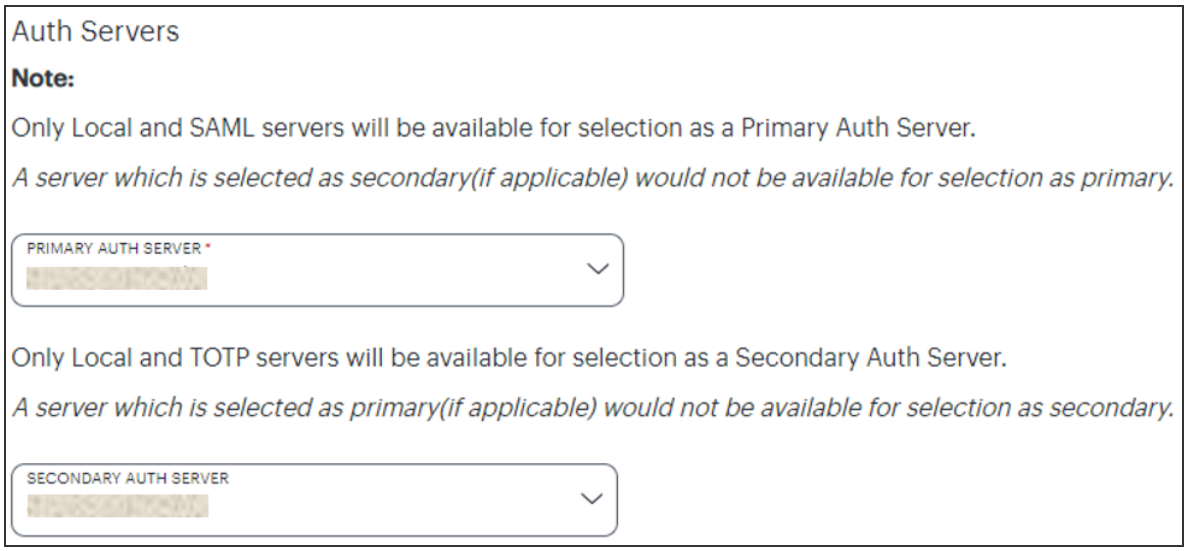

Selecting a primary authentication method for this policy

9. (Optional) Where a secondary method is required for Multi-Factor Authentication, select **Secondary Auth Server** from the drop-down list.

   Secondary authentication methods do not apply to *Enrollment User* type policies. The relevant field is hidden in this case.

10. Click **Create User Policy** to create the new policy.

    The new policy is added to the list of authentication policies.

If you instead elect to update an existing custom or built-in policy:

1. Click the three dots adjacent to the relevant policy, then select **Edit**.

   The **Edit authentication policy** form appears.

   For built-in authentication policies, all properties except **Primary Auth Server** and **Secondary Auth Server** (where applicable) are read-only.

2. Configure the primary and/or secondary authentication methods, as required:

   - Set the **Primary Auth Server** to be the new local user authentication method (indicated):

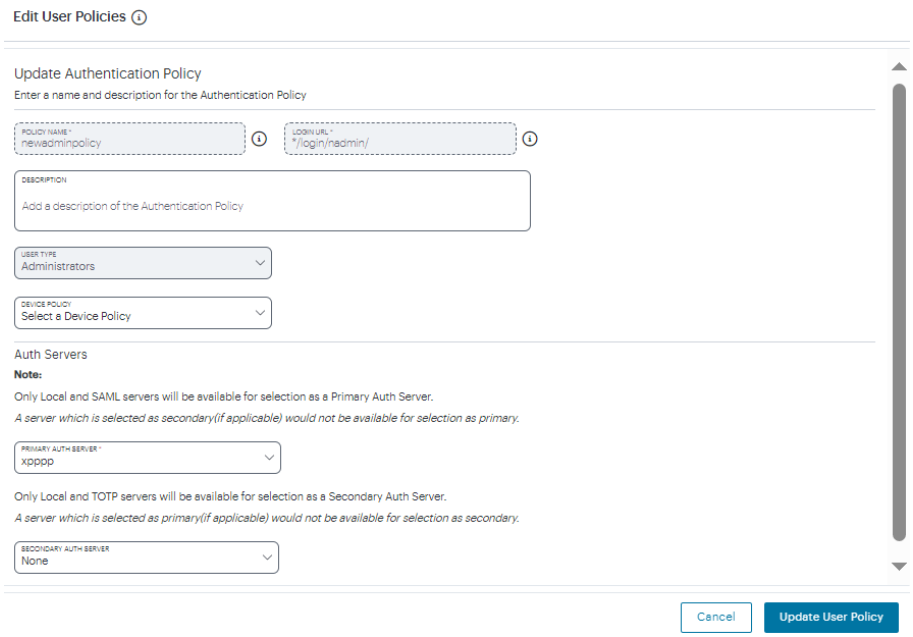

   Editing the primary auth server

   - If you are configuring a policy for MFA, set the **Secondary Auth Server** to be the new local user authentication method (indicated):

> If you configure a secondary authentication method in a policy that is currently in use, any active user sessions must be disconnected and reconnected for the changes to take effect.

3. Click **Update User Policy**.

   The list of authentication policies updates.

4. Repeat until all required authentication policies are updated.

To ensure that your users can access the authentication mechanism defined in the policies you configure through this process, make sure your *Secure Access Policies* are configured with a **User Group** in which these authentication policies are defined. To learn more, see Creating/Editing Secure Access Policies.

Workflow: Creating a SAML Authentication Policy With Azure AD

nZTA supports the use of a cloud-based Active Directory (AD) SAML service to provide authentication for your users.

If you choose to use AD as a SAML Identity Provider (IdP), you do not create any users locally on the *Controller*. All users will already be present in your remote SAML service.

Configuring *nZTA* to use SAML authentication requires you to create separate SAML apps on the Azure AD platform for the following primary activities:

- User sign-in
- User enrollment

As part of hardening custom sign-in policies and login URLs, the following changes are implemented:

1. Instead of requiring administrators to configure enrollment policies, administrators will only need to configure user policies. As a default, all configured user policies support enrollment.
2. Single SAML authentication server for user authentication and enrollment.

The *Controller* includes built-in default authentication policies for each of these purposes, and also includes the ability to create your own custom policies for separate authentication of specific user groups. You create an authentication method referencing one of the Azure AD SAML apps described above and then assign the method to an authentication policy of the same type (either the built-in policy, or one you create).

Configuring *nZTA* to Use SAML Authentication

You must keep the SAML metadata up-to-date, especially after renewing certificates. This is essential for a secure and successful SaaS Apps SAML SSO flow. Regularly updating configurations in both the Identity Provider and Service Providers helps prevent authentication failures and ensures the security of the authentication process.

Configure *nZTA* by performing the following steps:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > Authentication Servers**.

   The *Authentication Servers* page appears. This page lists all existing user authentication methods:

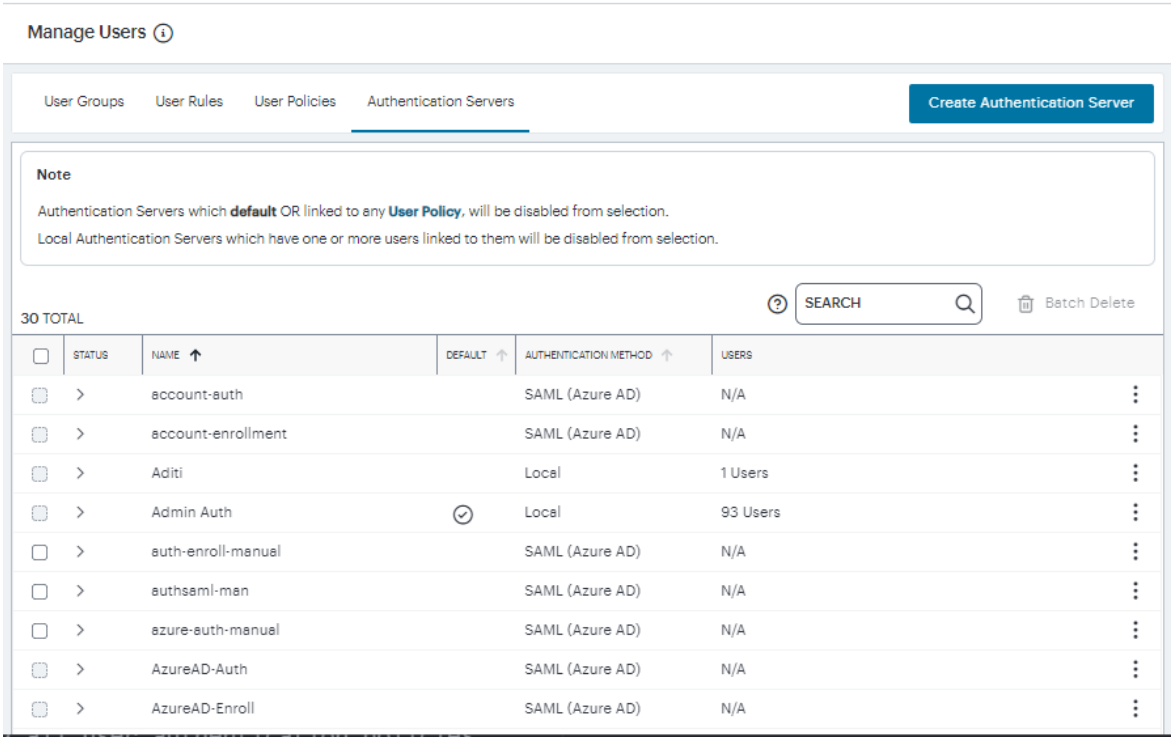

User Authentication Methods

3. Click **Create Authentication Server**.

   A form appears that enables you to define the authentication method:

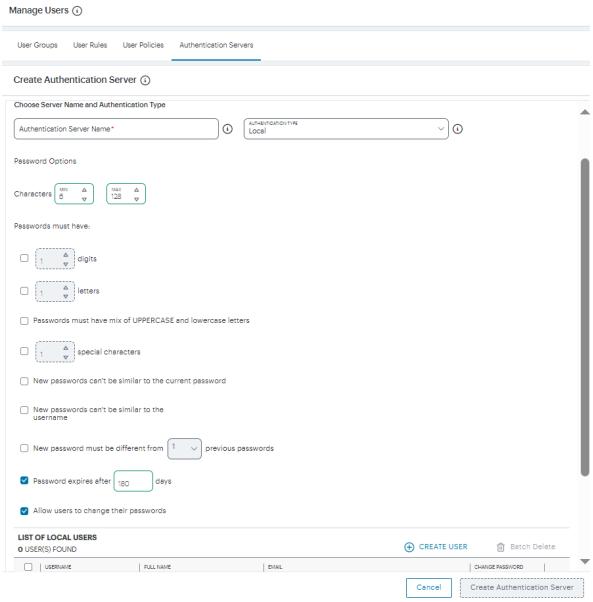

   Adding a user authentication method

> At any point during this process, you can reset the form data by selecting **Reset Fields**.

4. Under **Choose Server Name and Authentication Type**:

   - Select the **Authentication Type** as *SAML (Azure AD)*.

     The form expands to show additional settings.

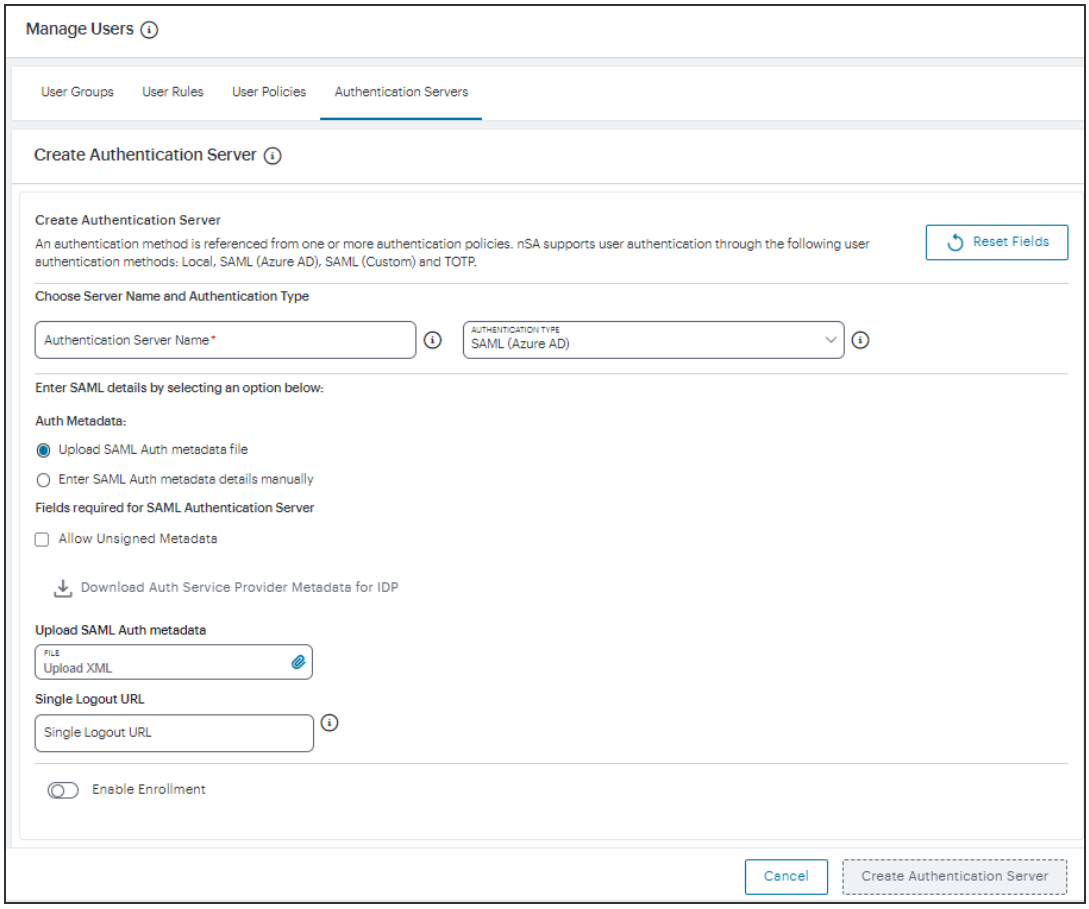

Configuring SAML (Azure AD) authentication settings

   - Specify an **Authentication Server Name**.

5. To provide your SAML IdP settings, select one of the following:

   - Select **Upload SAML Auth metadata file** if not selected already. This is selected by default.

     The **Download Auth Service Provider Metadata for IDP** link is enabled.

     1. (Optional) Specify a **Single Logout URL**. For more information, see Using SAML Single Logout to Force User Authentication.

        

        By default, the *Controller* expects a signed metadata definition file. To allow an unsigned metadata file, select **Allow Unsigned Metadata**.

     2. Click the **Download Auth Service Provider Metadata for IDP** link. Retain the downloaded file for later use.

- Select **Enter SAML Auth metada details manually** to manually enter the required IdP SAML settings. Use this option in scenarios where a SAML federation metadata file is not available or incomplete.

  Then, enter the following details:

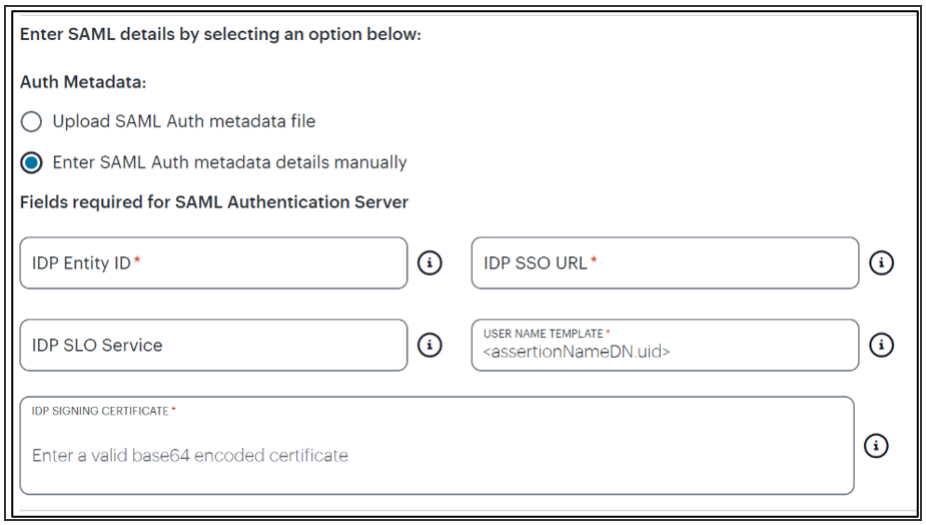

  Configuring SAML (Azure AD) IdP settings manually

  The following minimum settings are required for your SAML authentication service to function correctly. Each setting relates to a value configured in the SAML application on your IdP. Contact your IdP administrator to obtain the relevant details:

  - **IDP Entity ID**: The entity ID is sent as the Issuer value in the SAML assertion generated by the IdP. Specify the Issuer value in assertions generated by the SAML identity provider.
  - **IDP SSO URL**: A URL provisioned by the SAML IdP to support service-provider-initiated Single Sign-On. Use the format `https://<FQDN>`.
  - **IDP SLO Service**: (Optional) A URL to specify the Single Log-Out/sign out endpoint if you want to force re-authentication for increased security. Use the format `https://<FQDN>`. For more information, see Using SAML Single Logout to Force User Authentication.

- **User Name Template**: Specify how the system is to derive the username from the SAML assertion. The default value can be used, or replaced with an alternative specifier that the *Controller* uses from the incoming assertion. For example: `<assertionNameDN.uid>`, the NameID value where *ICS* is the IdP, the UID from X509SubjectName, `<userAttr.attr>`, attr from AttributeStatement attributes.
- **IDP Signing Certificate**: The signing certificate to be used with the SAML app on the IdP. Type or paste in the contents of your Base-64 encoded public key.

6. If this Auth server is used with User Policy of type "User", then click **Enable Enrollment**.

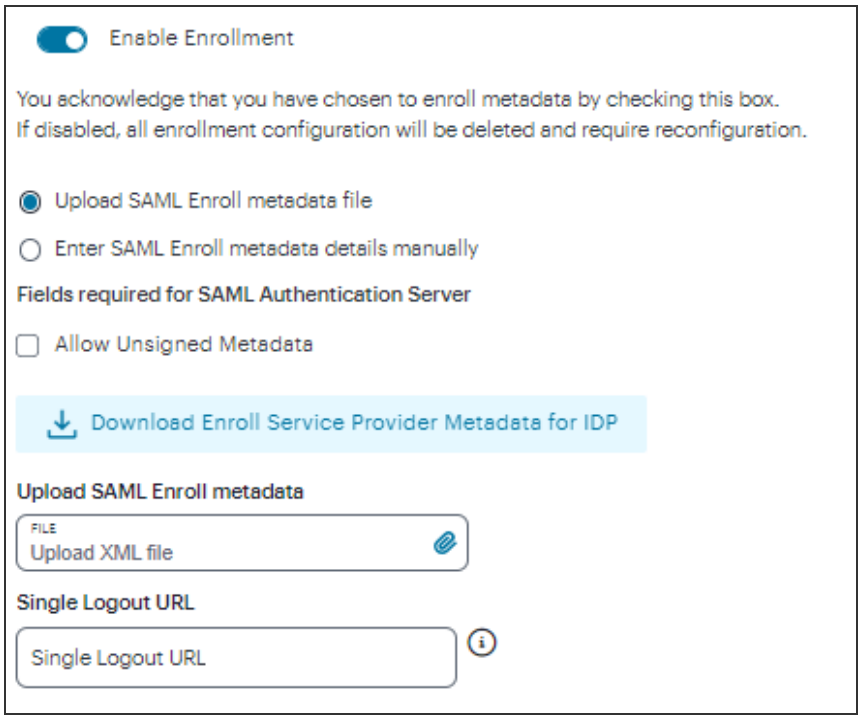

Enable Enrollment

7. To provide your SAML IdP settings, select one of the following:

    - Select **Upload SAML Enroll metadata file** if not selected already. This is selected by default.

        1. (Optional) Specify a **Single Logout URL**. For more information, see Using SAML Single Logout to Force User Authentication.

            By default, the *Controller* expects a signed metadata definition file. To allow an unsigned metadata file, select **Allow Unsigned Metadata**.

        2. Click the **Download Enroll Service Provider Metadata for IDP** link.

- Select **Enter SAML Enroll metada details manually** to manually enter the required IdP SAML settings. Use this option in scenarios where a SAML federation metadata file is not available or incomplete.

  Then, enter the following details:

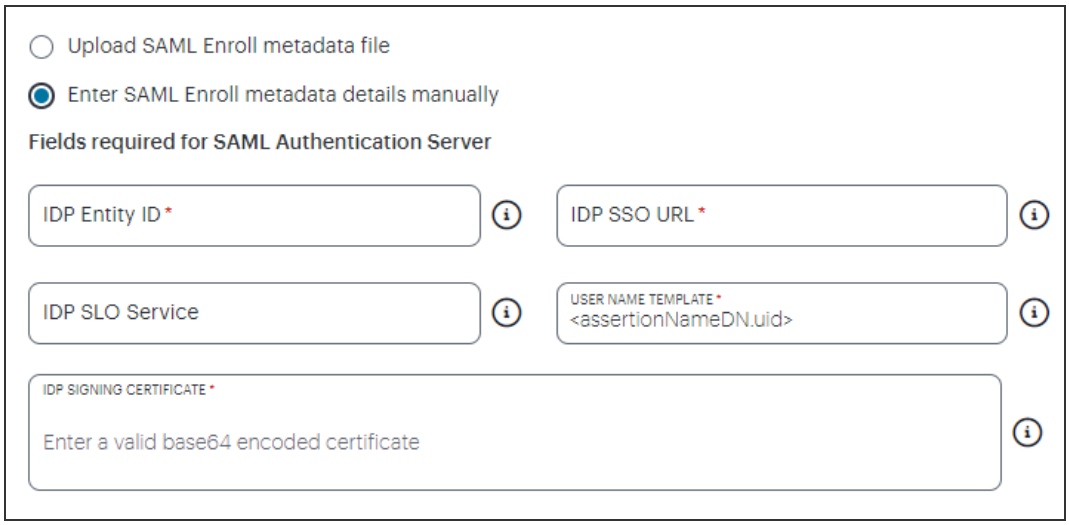

  Configuring SAML (Azure AD) IdP settings manually

  The following minimum settings are required for your SAML authentication service to function correctly. Each setting relates to a value configured in the SAML application on your IdP. Contact your IdP administrator to obtain the relevant details:

  - **IDP Entity ID**: The entity ID is sent as the Issuer value in the SAML assertion generated by the IdP. Specify the Issuer value in assertions generated by the SAML identity provider.
  - **IDP SSO URL**: A URL provisioned by the SAML IdP to support service-provider-initiated Single Sign-On. Use the format `https://<FQDN>`.
  - **IDP SLO Service**: (Optional) A URL to specify the Single Log-Out/sign out endpoint if you want to force re-authentication for increased security. Use the format `https://<FQDN>`. For more information, see Using SAML Single Logout to Force User Authentication.

- **User Name Template**: Specify how the system is to derive the username from the SAML assertion. The default value can be used, or replaced with an alternative specifier that the *Controller* uses from the incoming assertion. For example: `<assertionNameDN.uid>`, the NameID value where *ICS* is the IdP, the UID from X509SubjectName, `<userAttr.attr>`, attr from AttributeStatement attributes.
- **IDP Signing Certificate**: The signing certificate to be used with the SAML app on the IdP. Type or paste in the contents of your Base-64 encoded public key.

- When editing an existing SAML Auth server, the 'Enable Enrollment' option can be enabled or disabled if the SAML Auth server is not being used in any 'User Policy'. If the SAML Auth server is being used in a 'User Policy', then Enable Enrollment button will be grayed out.
- If 'Enable Enrollment' is not selected, then while creation of 'User Policy' of type 'User', the server you have created (without Enable Enrollment) will not be listed.
- When Enrollment is disabled, the enrollment SAML configuration will be deleted. To enable enrollment, you have to again provide enroll SAML auth server configuration.

8. Use the downloaded User Authentication and Enrollment SP Metadata files to create Sign-In and Enrollment SAML Apps in Azure AD. For details, see "Creating Enrollment/Sign-in SAML App in Azure AD" below.
9. Browse and upload the digitally-signed (or unsigned) federation metadata XML definition files downloaded from Azure AD.
10. Confirm that your settings are correct, then select **Create Authentication Server** to create the authentication method.

    The new SAML user authentication method is added to the list of methods displayed in the **User Authentication** page, and the process completes.
11. (Optional) To edit a listed authentication method, click the adjacent three dots, then select **Edit**. Make any required updates and confirm.
12. (Optional) To delete one (or more) an *unused* authentication methods, select the check box for each, then select **Delete**. You must confirm the deletion.

Creating Enrollment/Sign-in SAML App in Azure AD

Perform the following steps in Azure AD:

1. Create a SAML app for the required activity (sign-in or enrollment).
2. Click **Upload Metadata File** and select the file from your download.

   This defines at least the following **Basic SAML Configuration fields**:

   - **Identifier (Entity ID)**. This is the URL of the SAML endpoint on the *Controller*. This is the audience of the SAML response for IdP-initiated SSO. This cannot be left blank.
   - **Reply URL (Assertion Consumer Service URL)**. This is the URL of the SAML consumer on the *Controller*. This is the destination URL in the SAML response for IdP-initiated SSO. This cannot be left blank.

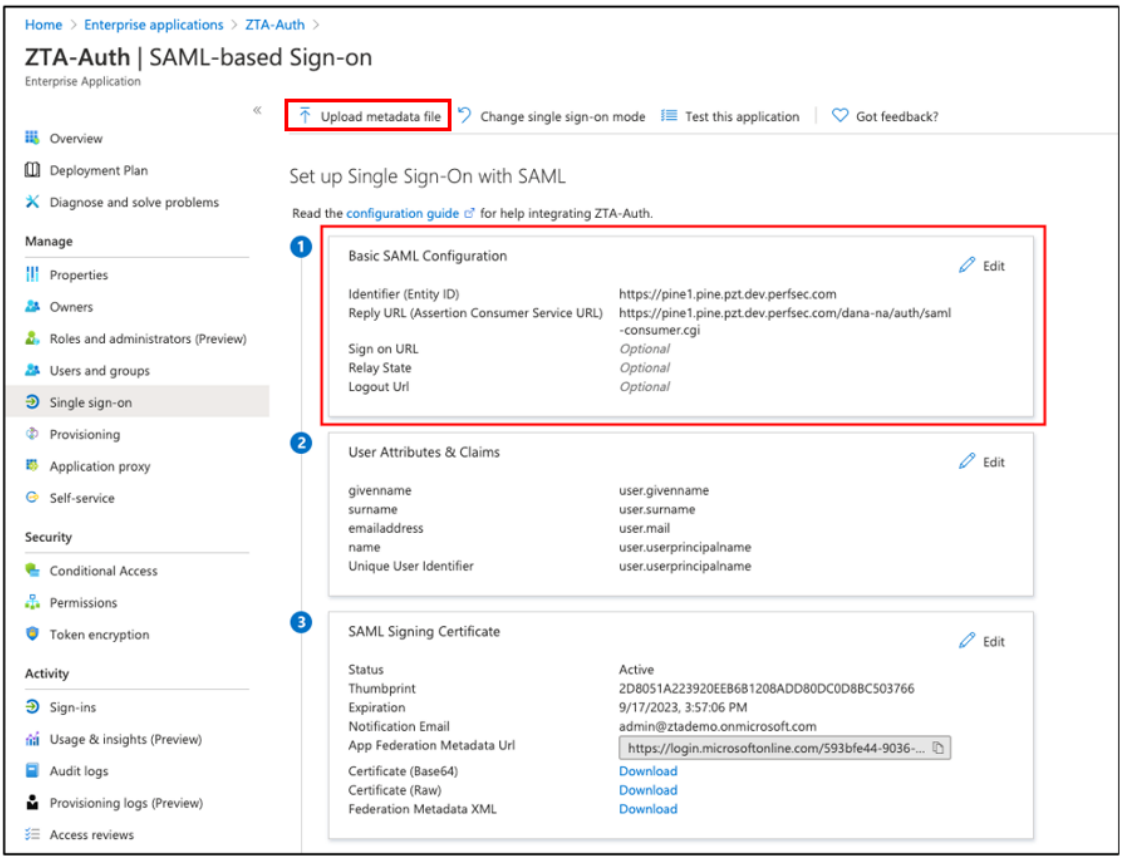

Setting *Basic SAML Configuration* in Azure AD applications

3. Download the Federation metadata XML definition for the SAML app to your local workstation. Retain this file for later use.

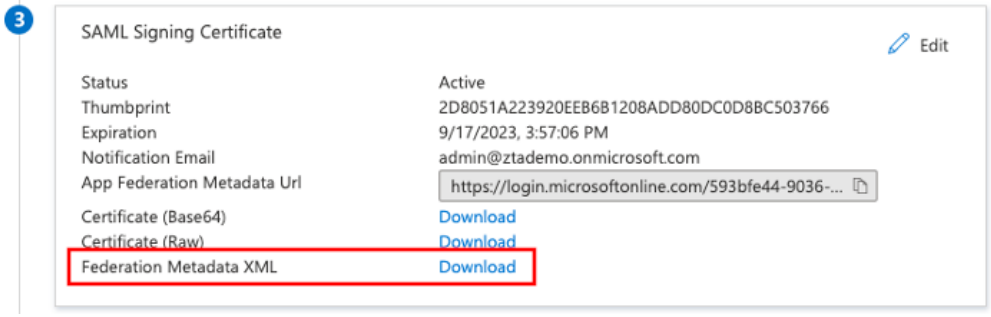

Downloading Federation Metadata XML files for user enrollment and user sign-in SAML applications

4. Repeat these steps for each activity.

> For details on how to create SAML apps in Azure AD, see the Azure AD SAML documentation.

Creating/Updating Authentication Policies

After you have created your SAML authentication method, create or update your authentication policies with the new authentication method.

1. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > User Policies**.

   The **User Policies** page appears. This page lists all existing user authentication policies.

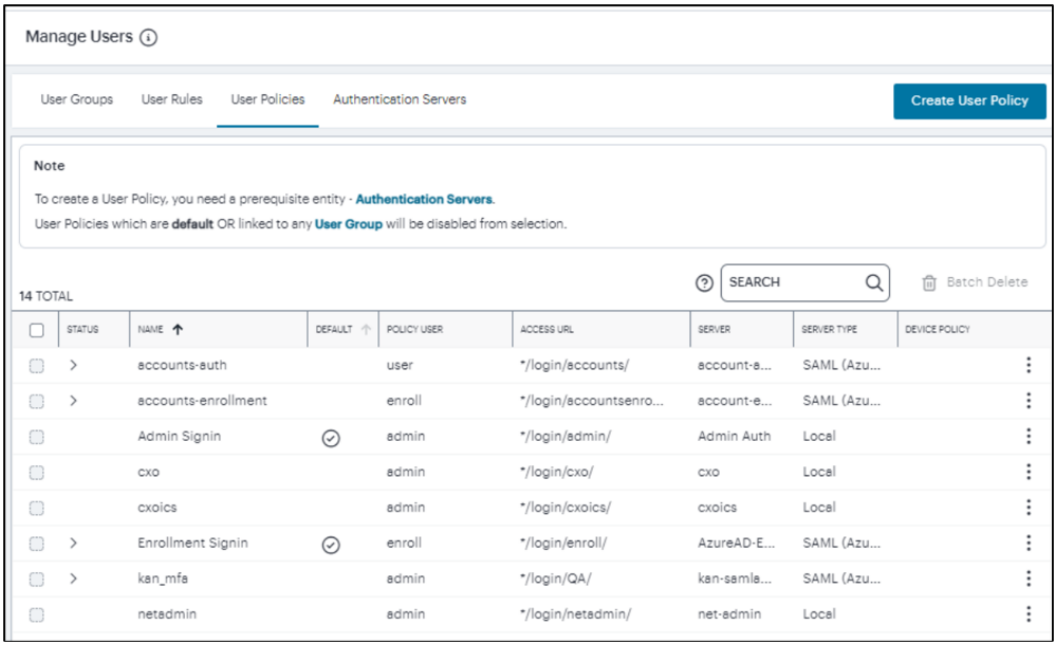

   User Authentication Policies

   To learn more about the policies on this page, see Viewing User Authentication Policies.

   From this page, either create a new custom policy or edit an existing policy.

2. To add a new custom policy, select **Create User Policy**.

   The **Create Authentication Policy** form appears.

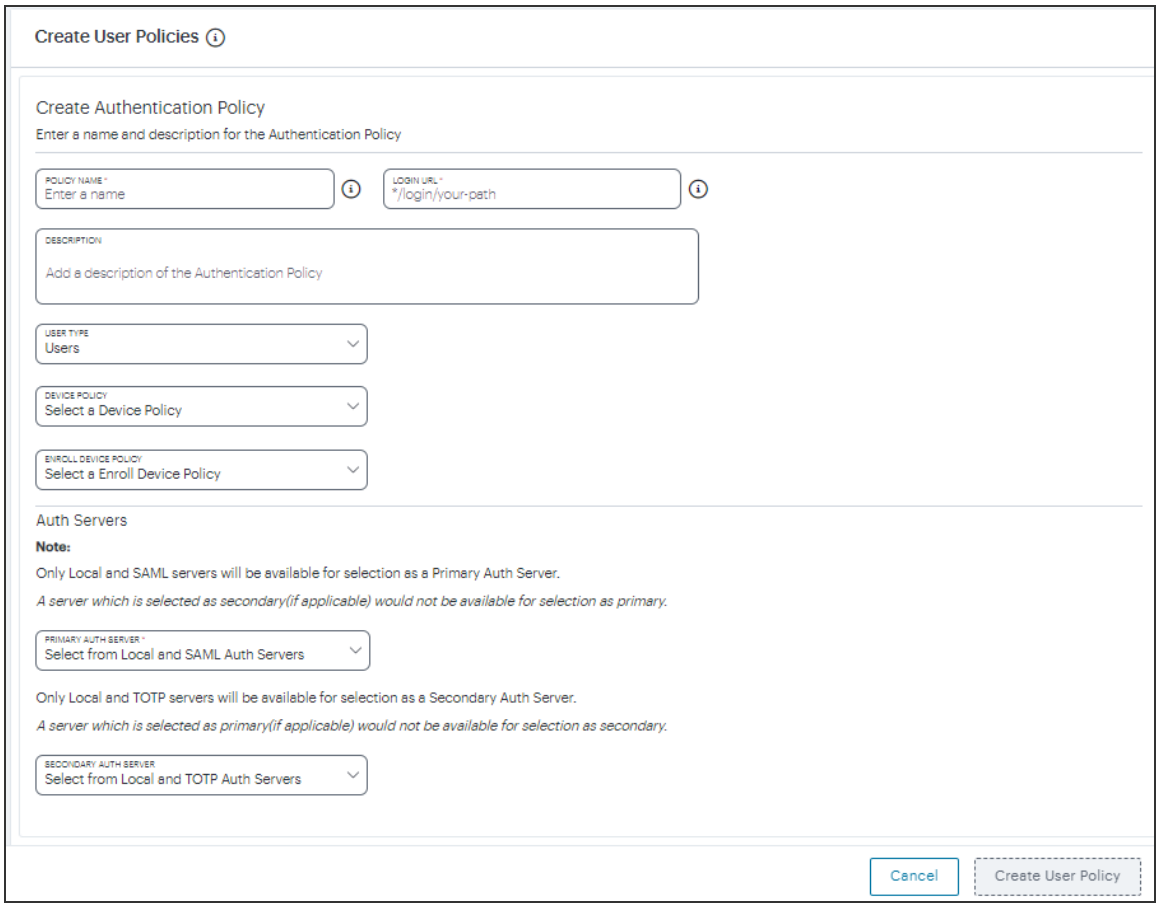

   Create Authentication Policy

> At any point during this process, you can reset the form data by selecting **Reset Fields**.

> To learn more about how custom policies are used for user login and enrollment, see Adding Custom Authentication Policies.

3. Enter a **Policy Name**.

4. Enter a **Login URL** using the format **/login/<path>/*.

   The URL must start with "**/login/*" and cannot contain any special characters. <path> should be set to a unique value reflecting the endpoint URL you want to define for this authentication policy (appended with a backslash):

   - In the case of user sign-in policies, this is the URL endpoint (appended to the tenant FQDN) to which new users are invited to connect to enroll or sign-in a device with the *Controller*. Example value: "**/login/saleslogin/*".
   - In the case of user enrollment policies, this endpoint identifies the enrollment URL to which users are redirected if they attempt to connect to the equivalent sign-in policy with an un-enrolled device. In most cases, you do not advertise this endpoint to your users. Example value: "**/login/salesenroll/*".

   

   In some enrollment circumstances, such as when using a device pre-installed with an older version of *Ivanti Secure Access Client*, you connect directly to the enrollment policy endpoint to enroll the device. For more details, see Using User Authentication Policies.

5. (Optional) Enter a description for the authentication policy.
6. Select a **User Type** based on the intended authentication activity for this policy. Choose from:

   - **Users**: Select this option to define the user sign-in endpoint for enrolled devices. This is the endpoint that you provide to your users to access the service (regardless of enrollment status).
   - **Administrators**: Select this option to define the authentication endpoint for administrator-level sign-in. This endpoint is used for administrator login to the *Controller* only.

7. From the **Device Policy** list, select a device policy.
8. (for policies with a **User Type** of "Users" only): Select an **Enroll Device Policy** from the drop-down list to be linked to this sign-in policy (as indicated):

   This is the enrollment policy to which a user is redirected if it is determined that the device is not yet enrolled. To learn more, see Using User Authentication Policies.

9. Under **Auth Servers**, select **Primary Auth Server** and choose the required authentication method from the drop-down list. Only Local and SAML servers will be available for selection as a Primary Auth Server.

10. (Optional) Where a secondary method is required for Multi-Factor Authentication, select **Secondary Auth Server** and choose the required authentication method from the drop-down list. Only Local and TOTP servers will be available for selection as a Secondary Auth Server.
11. Click **Create User Policy** to create the new policy.

    The new policy is added to the list of authentication policies.

Existing User Policies

- For existing default user policies, admin can select either existing or new SAML or Local Auth Server.

  Note: The Auth Server for enrollment will be added automatically.

- For existing custom user policies, user can update existing SAML Auth Server.

If you intend to update an existing custom or built-in policy:

1. Click the three dots adjacent to the relevant policy, then select **Edit**.

   The **Update Authentication Policy** form appears.

   

   For built-in authentication policies, all properties except **Primary Auth Server** and **Secondary Auth Server** (where applicable) are read-only.

2. Set the **Primary Auth Server** to be the new SAML user authentication method (indicated):

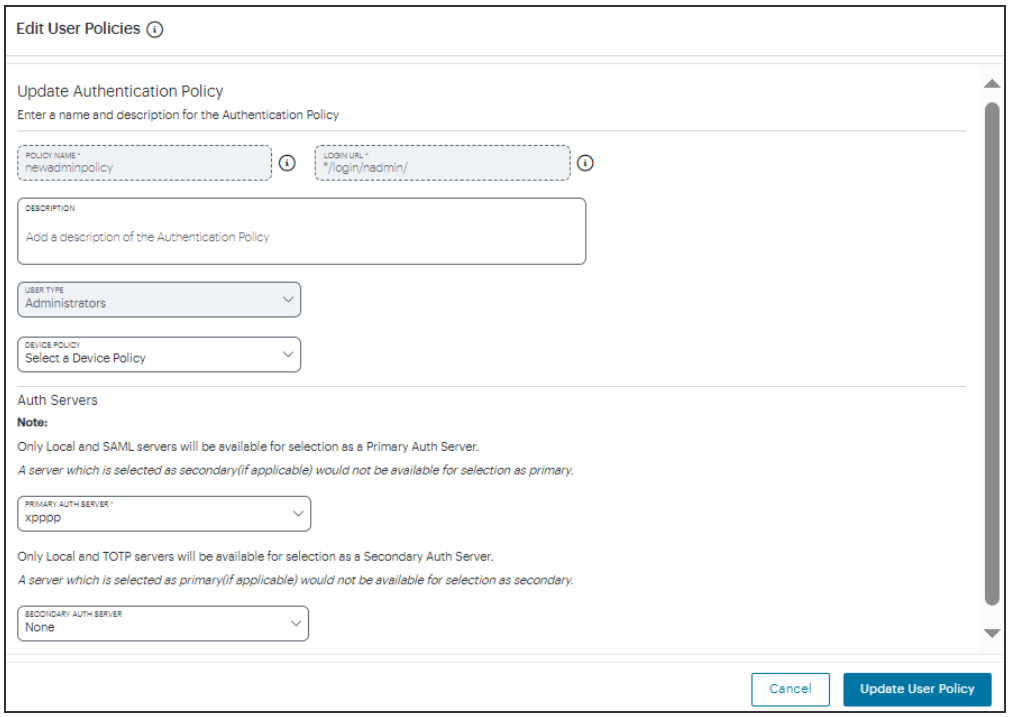

Editing the primary auth server

> SAML authentication can be used only as a Primary Auth Server. If you are using MFA, specify either a *local authentiction* or *TOTP* method as the **Secondary Auth Server**.

> If you configure a secondary authentication method in a policy that is currently in use, any active user sessions must be disconnected and reconnected for the changes to take effect.

3. Select **Update User Policy**.

   The list of authentication policies updates.

4. Repeat until all required authentication policies are updated.

At this point, the *Controller* uses the uploaded Federation Metadata to contact the SAML service. After this process completes, a **Download** function becomes available for each relevant policy. This metadata file is required to configure trusted communication with the remote SAML service.

1. Refresh your browser until the **Download** action is visible for the relevant policies.
2. Select the check box for the policy metadata you want to download and clear all other check boxes.
3. Select **Download** and save the metadata file.

> As mentioned previously, make sure you repeat this procedure for each required SAML app on your Azure AD platform. That is, you require separate XML metadata files for the enrollment authentication policy and the login authentication policy.

After the **User Authentication** workflow is complete, you can configure the Azure AD platform with the XML configuration of the *Controller*. For details on how to configure Azure AD, see Configuring Azure AD with Service Provider Metadata from the Controller.

Finally, to ensure that your users can access the authentication mechanism defined in the policies you configure through this process, make sure your *Secure Access Policies* are configured with a **User Group** in which these authentication policies are defined. To learn more, see Creating/Editing Secure Access Policies.

Configuring MFA with Azure AD (Optional)

> This section describes an optional configuration activity for SAML User Authentication on Azure AD.

Multi-Factor Authentication (MFA) is an authentication method by which a user is granted access only after successfully presenting two (or more) pieces of evidence (or factors) to an authentication mechanism.

The Azure AD platform supports MFA internally, and this can be configured so that two authentication factors are required by users for a single user authentication method.

After you have configured the Azure AD platform for SAML User Authentication, you can optionally add MFA support.

> For full details of this activity, consult the Azure AD documentation for *Azure AD Conditional Access* and *Multi-Factor Authentication*.

An overview of the general process is given below:

1. On Azure AD, locate the **Azure AD Conditional Access** page.
2. Create a **New Policy** with the following details:
    - Add a **Name** for the policy.
    - Under **Assignments > Users and Groups**:
        - Select **Select Users and Groups**.
        - Check the **Users and Groups** option, and select the required users and groups from the list.
        - Select **Done**.
    - Under **Assignments > Cloud Apps and Actions**:
        - Select **Select apps**.
        - Select (for example) the *zta-enroll* and *zta-auth* apps.
        - Select **Done**.
    - Under **Access Controls > Grant**:
        - Select **Require multi-factor authentication**.
    - Under **Enable Policy**:
        - Turn the policy **On**.
    - Select **Create**.
3. Locate the **Multi-Factor Authentication** page.
4. On the right-side panel, under **Configure**:
    - Select the **Additional cloud-based MFA settings**.

      A new tab appears.
    - Enable the following MFA methods as per the requirement.
        - *Call to phone*.
        - *Text message to phone*.
    - Select **Save**.

Configuration of MFA on Azure AD is complete.

Mobile end users will be asked for secondary authentication information from their next login.

You can then create user rules and user groups, see Creating User Rules and User Groups.

Configuring *nZTA* to use Azure AD Security Groups (Optional)

> This section describes an optional configuration activity for SAML User Authentication on Azure AD.

You can configure *nZTA* to retrieve Azure AD-based user security group information through SAML attributes applied to your *User Rules*.

You can create users and security groups directly in Azure AD portal. This means that Azure AD group information can be retrieved by the *Controller* through SAML attributes containing only the group's Object ID. To instead retrieve Azure AD group information via the `sAMAccountName` attribute, create an Azure AD group in the local Active Directory and synchronize it to the cloud using *Azure AD Connect*.

To create users and groups directly in the Azure AD portal:

> This process is applicable only if you are creating users or groups directly in Azure AD. For scenarios that involve synchronizing local Active Directory users or groups to the cloud, continue to the procedure that follows this.

1. Log in to the Azure AD portal
2. Navigate to your SAML SSO application, and select the **Users** page.

3. Add a new user for your organization by selecting **Create User**:

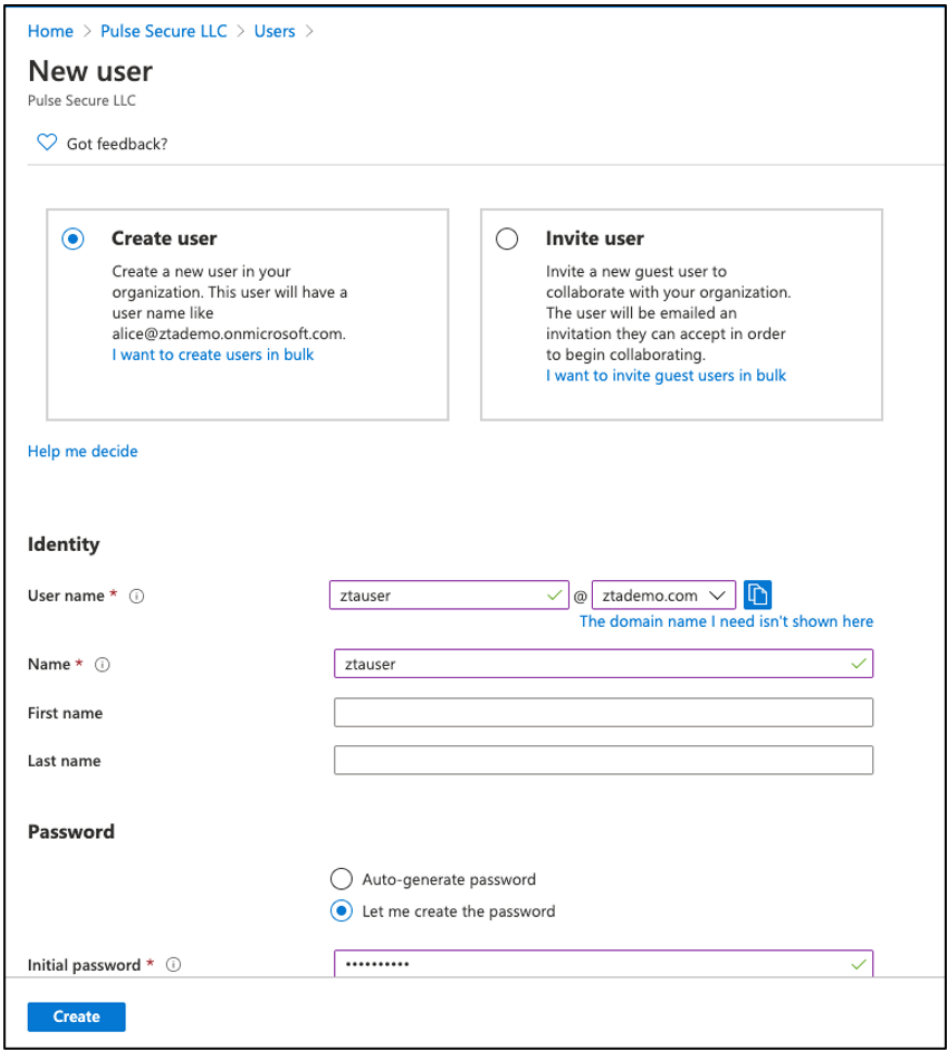

Adding a new user

Enter the following details:

- **User name**: Enter a username for the new user, appended by the domain applicable to your organization.
- **Name**: Enter the user's full display name.
- (Optional) **First Name**: Enter the user's first name.

- (Optional) **Last Name**: Enter the user's last name.
- **Password**: Select either an auto-generated or admin supplied password.

To create the user, select **Create**.

4. Add a new security group for your organization by navigating to the **Groups** page and selecting **New Group**:

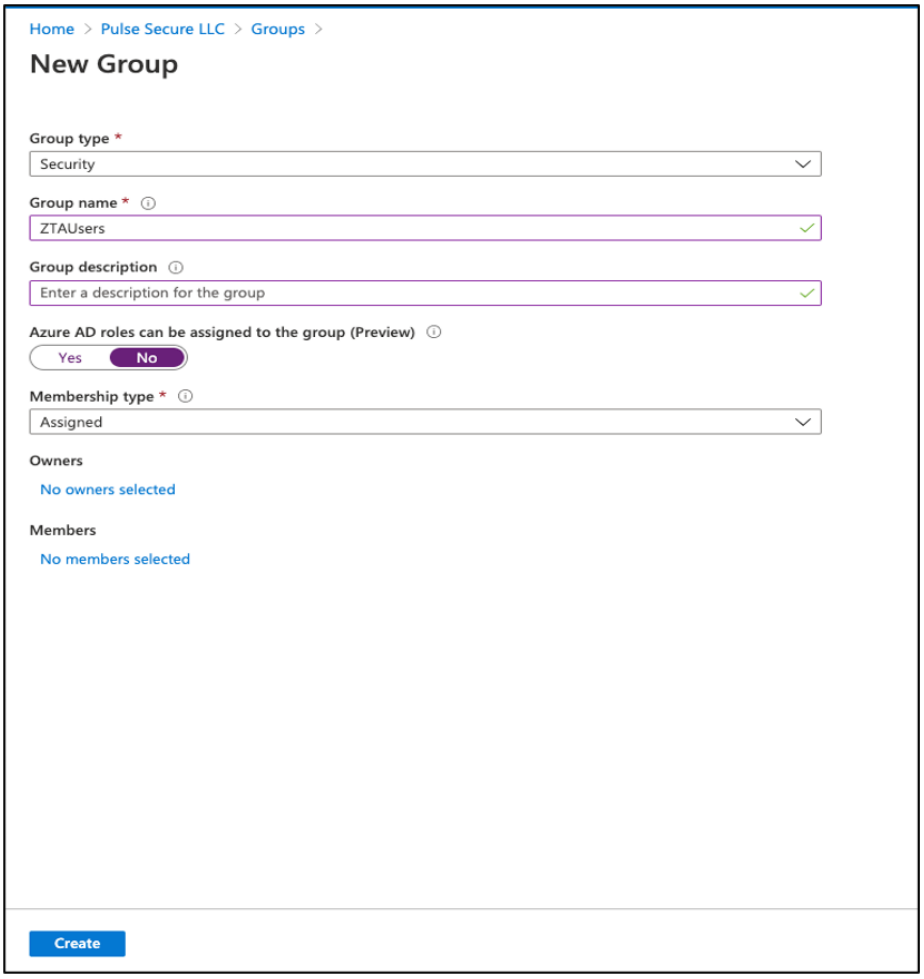

Adding a new security group

Enter the following details:

- **Group type**: Select *Security*.
- **Group name**: Enter a name for the new security group.
- (Optional) **Group description**: Enter a description for the group.
- (Optional) **Azure AD roles can be assigned to the group**: Use the default setting.
- **Membership type**: Select *Assigned*.

To create the group, select **Create**.

5. In the list of groups, make a note of the **Object ID** for the newly created group:

Viewing the **Object ID** for the new group

You use the **Object ID** property when configuring your *nZTAUser Rules*.

6. Assign the new user to your new security group:

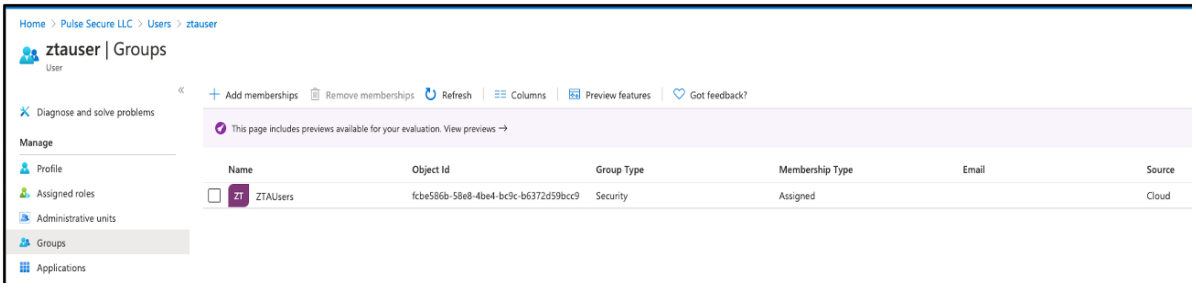

Assigning the new user to your new security group

To create Azure AD applications that facilitate retrieval of Security Group information:

1. Create two non-gallery applications under *Enterprise applications*, one for user enrollment and another for user sign-in, and configure Single Sign-On. Then, for each application, obtain the Federation Metadata XML file.

   For more details on this procedure, see Workflow: Creating a SAML Authentication Policy With Azure AD or follow the Azure AD user documentation.

2. Assign your applications to the required users or security groups:

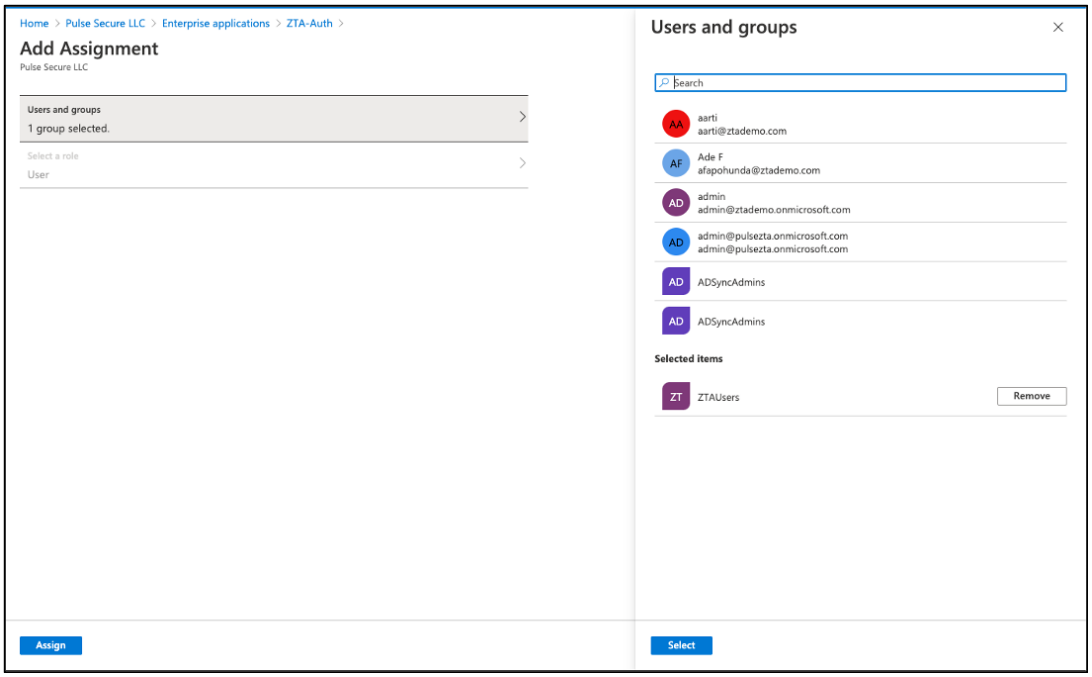

Assigning applications to users and security groups

3. Create two new *SAML (Azure AD)* user authentication methods (user enrollment and user sign-in) in the *nZTA* Tenant Admin portal that correspond to the SAML applications created in the Azure AD portal. Upload the Federation Metadata XML file from each application to the matching method.

4. Create corresponding authentication policies and assign each Azure AD authentication method as the **Primary Auth Server**.

5. Obtain the IdP metadata files for both authentication policies and upload this data to the Azure AD portal under the respective applications. For more details, see Configuring Azure AD with Service Provider Metadata from the Controller.

6. For each Azure AD application, in the *User Attributes and Claims* section, select **Edit**.

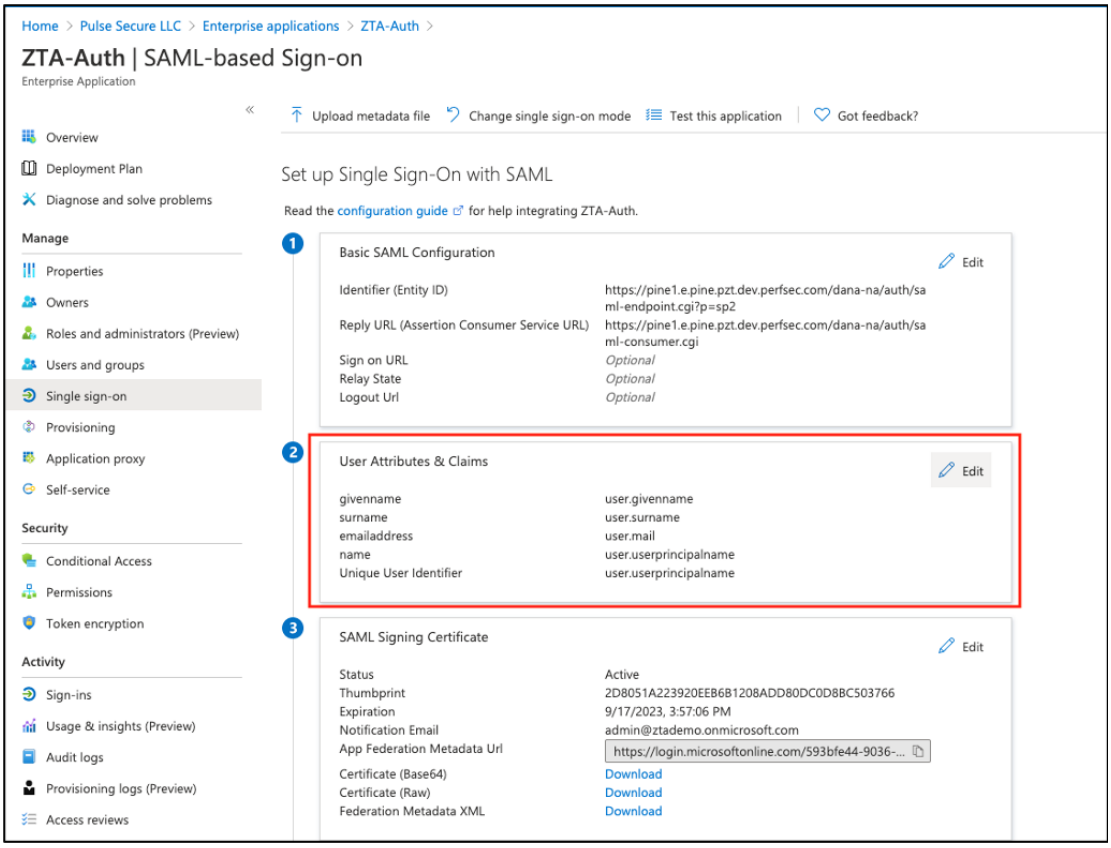

Editing the user attributes and claims for an Azure AD application

7. Add a new group claim to your application by selecting **Add a group claim**:

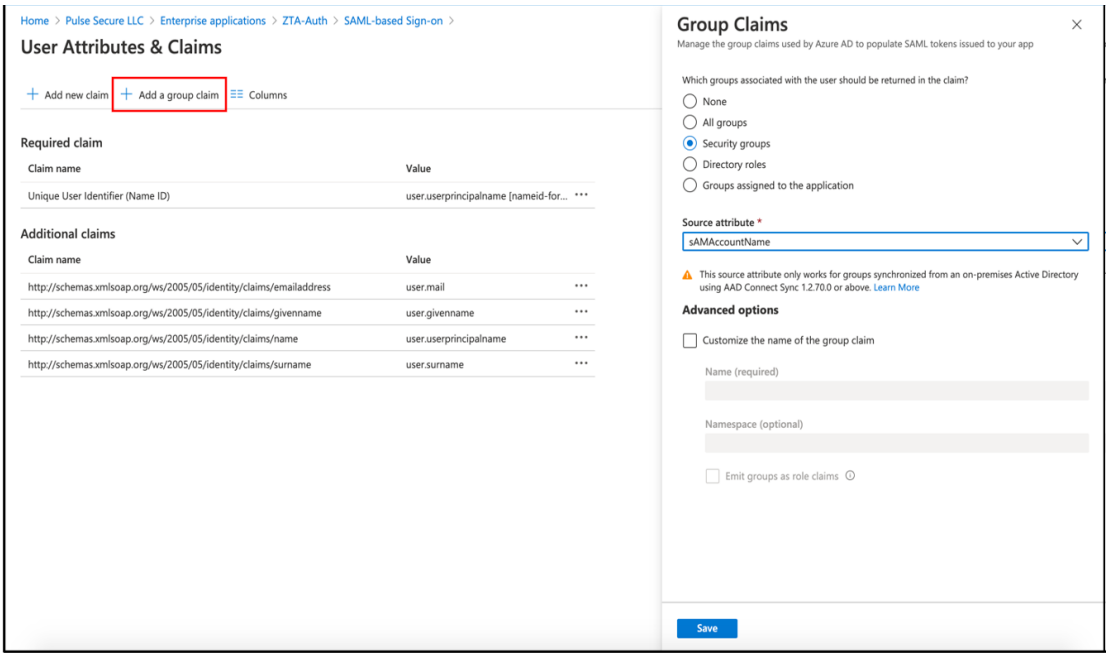

Adding a new group claim to your Azure AD application

In the *Group Claims* panel, enter the following details:

- **Which groups associated with the user should be returned in the claim?**: Select "Security groups".

- **Source attribute**:

    - To retrieve Azure AD security group details based on the group object ID, select "Group ID".
    - To retrieve Azure AD security group details based on the group name, select "sAMAccountName". This option is available only for groups synchronized from Active Directory.

    By default, Azure AD populates SAML attributes for group claims in the format:

    ```
    userAttr.http://schemas.microsoft.com/ws/2008/06/identity/claims/groups

    samlMultiVarAttr.http://schemas.microsoft.com/ws/2008/06/identity/claims/groups
    ```

    To specify a different claim type for group claims, select **Customize the name of the group claim** and specify the claim type in the **Name** field. For example, if you specify "ADgroup", the group claim is emitted as `userAttr.ADgroup` or `samlMultiValAttr.ADgroup`.

To add the new group claim, select **Save**.

8. Return to the *nZTA* Tenant Admin portal and navigate to the **Secure Access** > **Manage Users** > **User Groups and Rules** > **User Rules** page. Create a new user rule to access the Azure AD group attribute.

   Configure the following settings:

   - **Rule name**: Enter a descriptive name for the rule
   - **Select Attribute Type**: Select "SAML (Azure AD)".
   - **SAML Attribute Type**: Select "Group".

- **Attribute value**: Use one of the following options:
  - If you created the security group in Azure AD, use the group *Object ID* as noted earlier in the process:

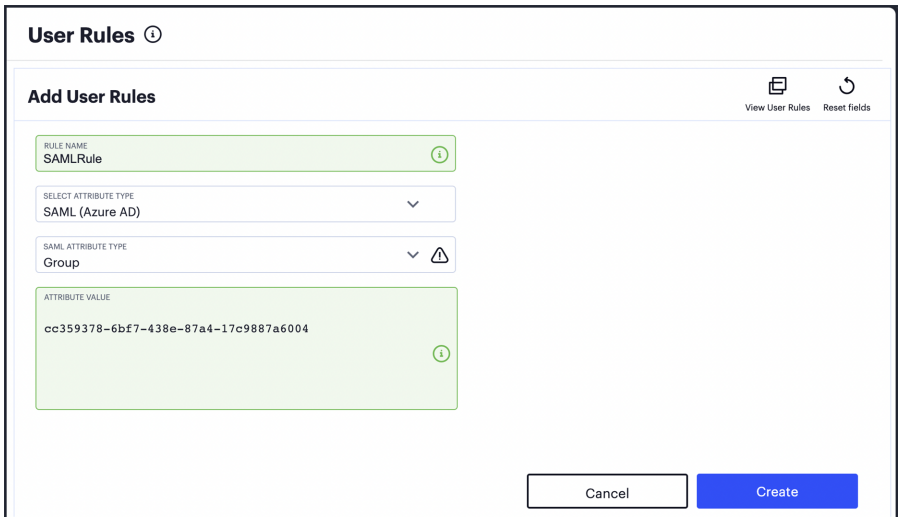

Specifying an Object ID as the group attribute value

  - If you synchronized your security group to Azure AD from a local Active Directory, use the *Group Name*.

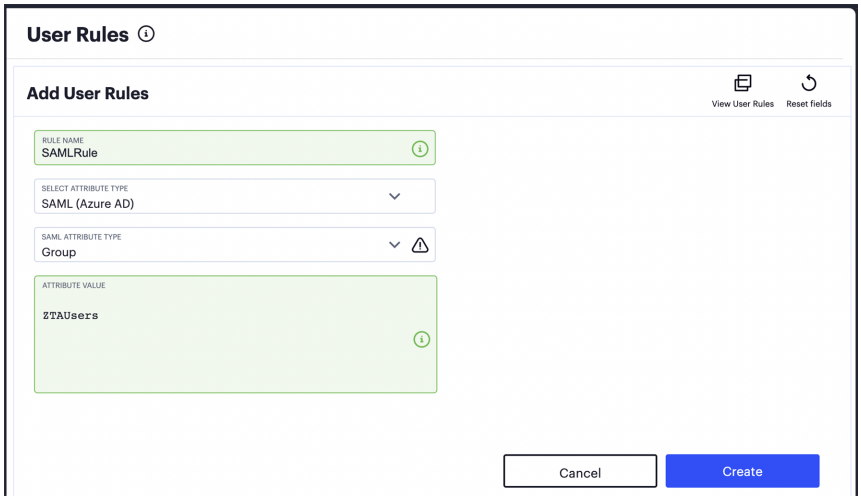

Specifying an Object Name as the group attribute value

The following expressions are examples of Azure AD group attribute values:

- Users that are members of only the identified group:

```
 userAttr.
{http://schemas.microsoft.com/ws/2008/06/identity/claims/groups}
= "<Group Object ID>"
```

- Users that are members of only the named group (for AD groups synchronized from a local Active Directory):

```
 userAttr.
{http://schemas\.microsoft\.com/ws/2008/06/identity/claims/group
s} = "<Group Name>"
```

- Users that are members of only the identified group:

```
userAttr.<Customized Group Claim Name> = "<Group Object ID>"
```

- Users that are members of only the named group (for AD groups synchronized from a local Active Directory):

```
userAttr.<Customized Group Claim Name> = "<Group Name>"
```

To add SAML multi-valued attributes (`samlMultiValAttr.<claim>`) in Azure AD User Rules, you must set **Select Attribute Type** to "SAML (Custom)". This applies to the following examples:

- Users that are members of either identified group:

```
samlMultiValAttr.
{http://schemas\.microsoft\.com/ws/2008/06/identity/claims/group
s} = ("<Group OID A>" or "<Group OID B>")
```

- Users that are members of both identified groups:

```
samlMultiValAttr.
{http://schemas\.microsoft\.com/ws/2008/06/identity/claims/group
s} = ("<Group OID A>" and "<Group OID B>")
```

- Users that are members of either named group (for AD groups synchronized from a local Active Directory):

```
 samlMultiValAttr.
{http://schemas\.microsoft\.com/ws/2008/06/identity/claims/group
s} = ("<Group A>" or "<Group B>")
```

- Users that are members of both named groups (for AD groups synchronized from a local Active Directory):

```
 samlMultiValAttr.
{http://schemas\.microsoft\.com/ws/2008/06/identity/claims/group
s} = ("<Group A>" and "<Group B>")
```

- Users that are members of either named group (for AD groups synchronized from a local Active Directory):

```
 samlMultiValAttr.<Customized Group Claim Name> = ("<Group A>" or
"<Group B>")
```

- Users that are members of both named groups (for AD groups synchronized from a local Active Directory):

```
 samlMultiValAttr.<Customized Group Claim Name> = ("<Group A>"
and "<Group B>")
```

9. To create the User Rule, select **Create**.

10. Create a new **User Rule Group** and select the User Authentication Policy (Enrollment or Sign-in) that is applicable to this service. Then, in the list of *User Rules*, select the rule you created in the previous step:

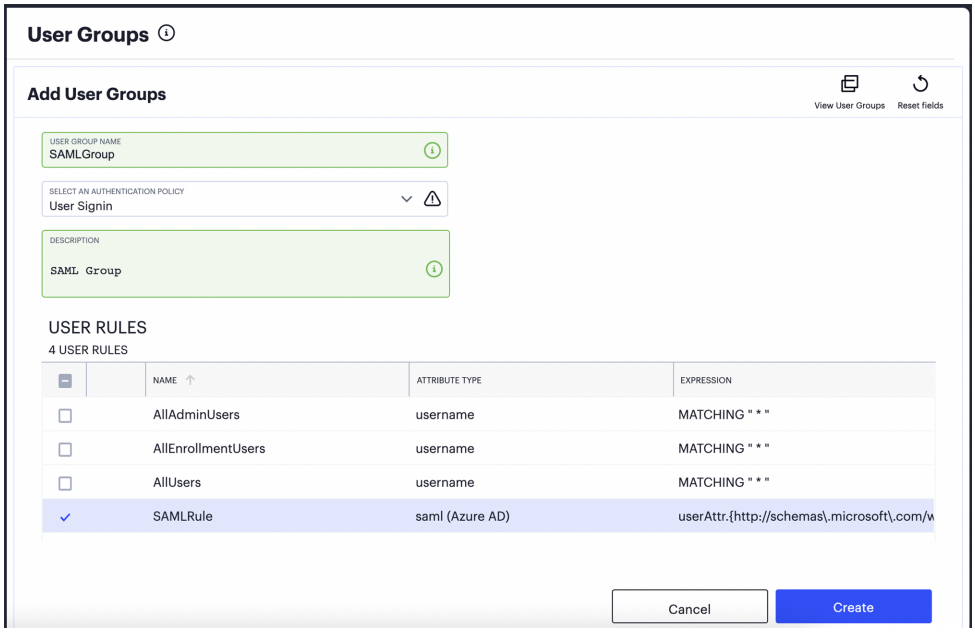

Configuring a User Rule Group with a User Authentication Policy and User Rule

11. To create the User Rule Group, select **Create**.

12. Proceed to create your **Secure Access Policy**, selecting the applicable *Application*, *Gateway*, *Device Policy* and the *User Rule Group* created during this procedure.

Workflow: Creating an Authentication Policy for On-Premises *ICS* SAML

You can choose to use a configured *ICS* server (either remote or local) as an on-premises SAML AD authentication server for your *Controller*.

> If you choose to use SAML authentication on your *Controller*, you do not create any users manually. All users will already be present on your remote SAML server.

Configuring *nZTA* to use SAML authentication requires you to create separate SAML apps on the on-premises *ICS* server for the following primary activities:

- User enrollment
- User sign-in

The *Controller* includes built-in default authentication policies for each of these purposes and includes the ability to create your own custom policies for separate authentication of specific groups. You create an authentication method referencing one of the SAML apps described above and then assign the method to an authentication policy of the same type (either the built-in policy, or one you create). Begin with enrollment, and then repeat the process for user sign-in.

To configure *nZTA* to use SAML authentication through *ICS*, perform the following steps:

1. Configure your on-premises *ICS* to act as a SAML AD authentication server:
    - Optionally, Configuring Secondary Authentication for On-Premises ICS (Optional).
    - Configuring a SAML Identity Provider in Ivanti Connect Secure.
    - Configuring a Metadata Provider in Ivanti Connect Secure.
2. Define an on-premises SAML authentication method, see Defining an On-Premises SAML Authentication Method.
3. Configure an on-premises SAML authentication policy, see Defining Authentication Policies for On-Premises SAML Authentication.
4. After you complete the User Authentication workflow, configure the SAML apps on the on-premises *ICS* server with the XML metadata configuration from the matching *nZTA* authentication policy, see Configuring ICS with Controller Metadata.

    You will need to repeat this process for each required SAML app on your *ICS* server.

To ensure that your users can access the authentication mechanism defined in the policies you configure through this process, make sure your *Secure Access Policies* are configured with a **User Group** in which these authentication policies are defined. To learn more, see Creating/Editing Secure Access Policies.

Configuring Secondary Authentication for On-Premises *ICS* (Optional)

This section describes an optional configuration activity for SAML User Authentication on an on-premises *ICS* server.

Multi-Factor Authentication (MFA) is an authentication method by which a computer user is granted access only after successfully presenting two (or more) pieces of evidence (or factors) to an authentication mechanism. The ICS platform supports a number of secondary authentication solutions, and can be configured so that two authentication factors are required by users using a single nZTA user authentication method. Before you configure the on-premises ICS server and Controller for SAML User Authentication, you can choose to configure MFA (secondary) support to ICS.

Before you start this procedure, you must have fully configured a secondary authentication server. For example, *Google OTP (One Time Password)* or *NAS OTP*. For the purposes of this example, an existing local *Google OTP* server is used, but different established secondary authentication methods are also supported.

To configure *ICS* for secondary authentication:

1. Log in to the on-premises *Ivanti Connect Secure* server.
2. On the **Authentication** menu, select **Auth. Servers**.

   The **Authentication Servers** page appears.
3. Under **New**, select the required authentication server type. For example, select *Time based One Time Password (TOTP) Server*.
4. Select **New Server**.

   The **New Time based One Time Password (TOTP) Server** page appears.
5. Enter the following parameters for the new server:
   - Add a **Name** for the server, for example *Google OTP*.
   - Select the **Allow Auto Unlock** check box, and set **Auto unlock period**. For example, 10 minutes.
   - Select the **Allow new TOTP user registration to happen via External Port** check box.
   - Select the **Allow TOTP authentication from Remote Pulse Secure Devices** check box.
   - Leave all other parameters as their default settings.
6. Select **Save Changes**.

   The new server (Google OTP) is added to the list of **Authentication Servers**.

7. You must now add a new user to the local system. To do this:
    - In the **Authentication Servers** list, select the hyperlink for the *System Local* entry.
    - The **Settings** page for *System Local* appears.
    - Select the **Users** tab.
    - Select **New** to create a user.
    - The **New Local User** page appears.
    - Enter details for the user, and ensure that the **Enabled** check box is selected, and that other check boxes are clear.
    - Select **Save Changes** to create the user.
8. Select the **Users** menu, then select **User Realms**.

    The **User Authentication Realms** page appears.
9. You must now enable secondary authentication. To do this:
    - In the **User Authentication Realms** list, select the hyperlink for the *Users* entry.

        The **General** page for the *Users* realm appears.
    - Under **Additional authentication server**, select the **Enable additional authentication server** check box.
    - Set **Authentication #2** to *Google OTP*.
    - At the bottom of the page, select **Save Changes**.

You can now configure an Identity Provider in *ICS*, see Configuring a SAML Identity Provider in Ivanti Connect Secure.

Configuring a SAML Identity Provider in *Ivanti Connect Secure*

This section describes the steps to configure a SAML Identity Provider (IdP) on an on-premises *Ivanti Connect Secure* (*ICS*) server. The metadata for the IdP is required by each SAML user method on *nZTA*.

To configure SAML IdP on the *Ivanti Connect Secure* server:

1. Log in to the on-premises *Ivanti Connect Secure* server that is identified as an Identity Provider.
2. Navigate to **System > Configuration > SAML > Settings**.
3. Configure the following Metadata Server Configuration:
    - **Timeout value for metadata fetch request** to 300.
    - **Host FQDN for SAML** to the Fully Qualified Domain Name, noting the host FQDN guidance below.

    The host FQDN specified here is used in the SAML entity ID, used by browsers to connect to *ICS*, and used in the URLs for SAML services. Typically:

    - If the *ICS* is standalone, the FQDN should resolve to the IP address of the external interface / internal interface, whichever is chosen.
    - If the *ICS* is an Active-Passive cluster, the FQDN should resolve to the external VIP / Internal VIP, whichever is chosen.
    - If the *ICS* is an Active-Active cluster behind an in-line load balancer, the FQDN should resolve to the load balancer's external VIP / Internal VIP, whichever is chosen.
4. Select **Save Changes**.
5. Select **Update Entity IDs** and confirm this action on the warning page by selecting **Update Entity IDs**.
6. Navigate to **System > Configuration > Certificates > Device Certificate**, create a new CSR, and import certificate and keys. Skip this step if the *ICS* external interface / internal interface (whichever is chosen) already provides a certificate that matches the host's Fully Qualified Domain Name.
7. Navigate to **Authentication > Signing In > Sign In SAML > Identity Provider**.
8. Locate the the **Basic Identity Provider (IdP) Configuration** section.
9. Under **Protocol Binding to use for SAML Response**:
    - Select the **Post** check box.
    - Clear the **Artifact** check box.
    - Select the required **Signing Certificate**.

10. Under Other Configurations:
    - Select the **Accept unsigned AuthnRequest** check box.
    - Select **sha-256** for Signature Algorithm, if ICS is using 22.x version onwards.
11. Under Service-Provider-related IDP Configuration:
    - For **SignIn Policy**, select the */* policy.
12. Under **User Identity**:
    - For **Subject Name Format**, select Email Address.
    - For **Subject Name**, enter <USERNAME>.
    - At the bottom of the page, select **Save Changes**.

You can now configure a Metadata Provider in *ICS*, see Configuring a Metadata Provider in Ivanti Connect Secure.

Configuring a Metadata Provider in *Ivanti Connect Secure*

This section describes the steps to configure *Ivanti Connect Secure* to be a Metadata Provider, and to download metadata for use on the *Controller*.

To configure a Metadata Provider in the on-premises *ICS* server:

1. Log in to *Ivanti Connect Secure* server.
2. Navigate to **Authentication > Signing-In > Sign In SAML > Metadata Provider**.

   The SAML Metadata Provider Entity Id property is pre-populated. It is generated by the system, based on the value for the Host FQDN for SAML setting on the **System > Configuration > SAML > Settings** page.
3. Set **Metadata Validity** to 365 days.
4. Clear the **Do Not Publish IdP in Metadata** check box.
5. Select **Save Metadata Provider**.
6. Select **Download Metadata** and save the file to your computer.

   This definition file is required to enable on-premises SAML apps on the *Controller*.

You can then configure the *Controller* to use an on-premises SAML Server, see Defining an On-Premises SAML Authentication Method.

Defining an On-Premises SAML Authentication Method

A minimum of two authentication methods are required:

- An authentication method for a SAML enrollment sign-in app.
- An authentication method for a SAML user sign-in app.

To create an on-premises SAML server authentication method for a specific activity, for example, device enrollment:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > Authentication Servers**.

   The *Authentication Servers* page appears. This page lists all existing user authentication methods:

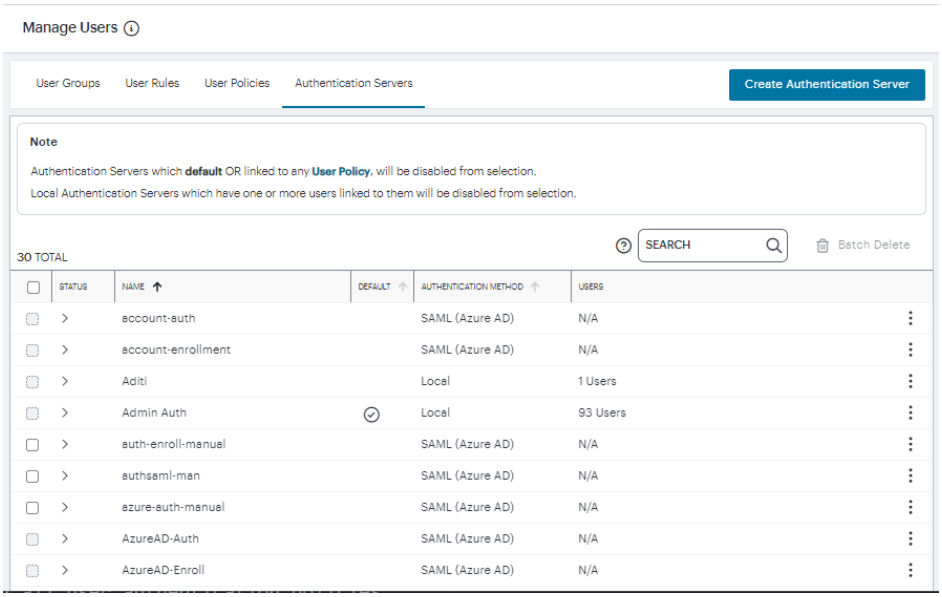

User Authentication Methods

3. Select **Create Authentication Server**.

   A form appears that enables you to define the authentication method:

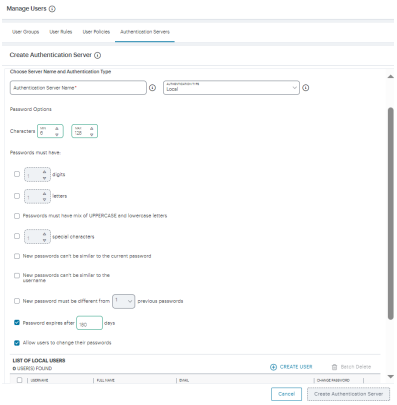

   Creating a user authentication method

> At any point during this process, you can reset the form data by selecting **Reset Vields**.

4. Under **Choose Server Name and Authentication Type**:

   - Select the **Authentication Type** of *SAML (Custom)*.

     The form expands to show additional settings:

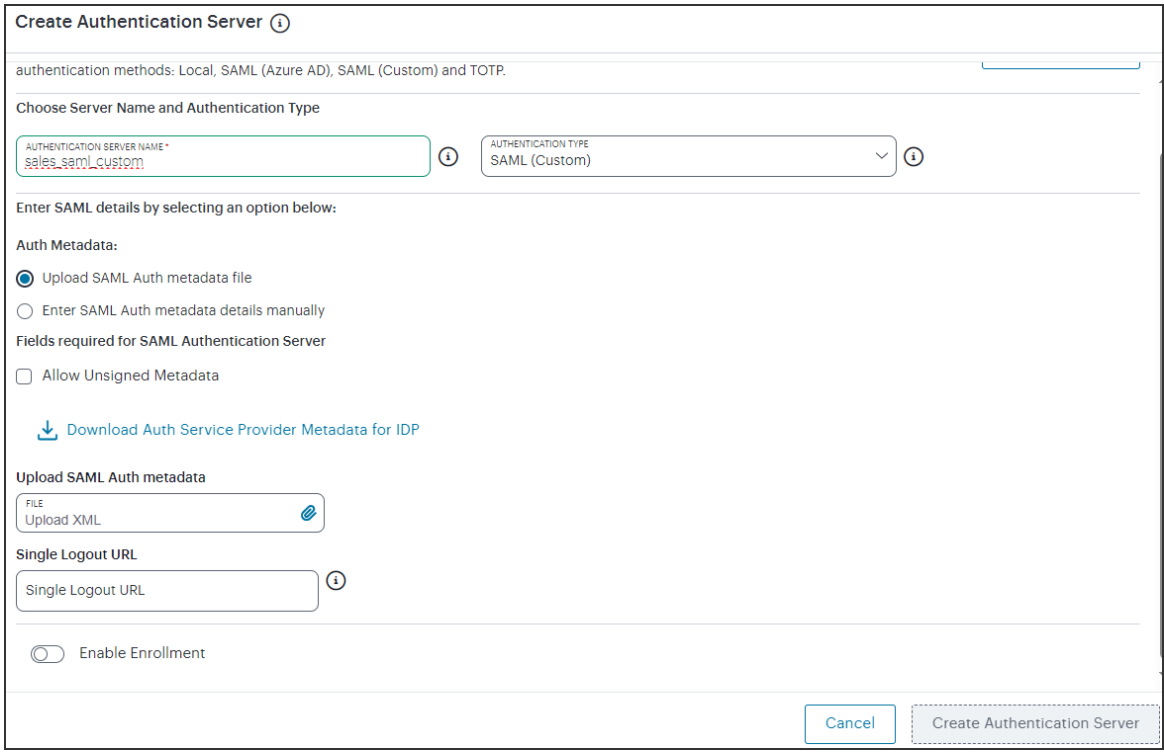

     Configuring SAML (Custom) authentication settings

   - Specify an **Authentication Server Name**. For example: *Enrollment* or *SignIn*.

5. To provide your SAML IdP settings, select one of the following:

   - Select **Upload SAML Auth metadata file** if not selected already. This is selected by default.

     The **Download Auth Service Provider Metadata for IDP** link is enabled.

     1. (Optional) Specify a **Single Logout URL**. For more information, see Using SAML Single Logout to Force User Authentication.

        

        By default, the *Controller* expects a signed metadata definition file. To allow an unsigned metadata file, select **Allow Unsigned Metadata**.

     2. Click the **Download Auth Service Provider Metadata for IDP** link. Retain the downloaded file for later use.

- Select **Enter SAML Auth metada details manually** to manually enter the required IdP SAML settings. Use this option in scenarios where a SAML federation metadata file is not available or incomplete.

  Then, enter the following details:

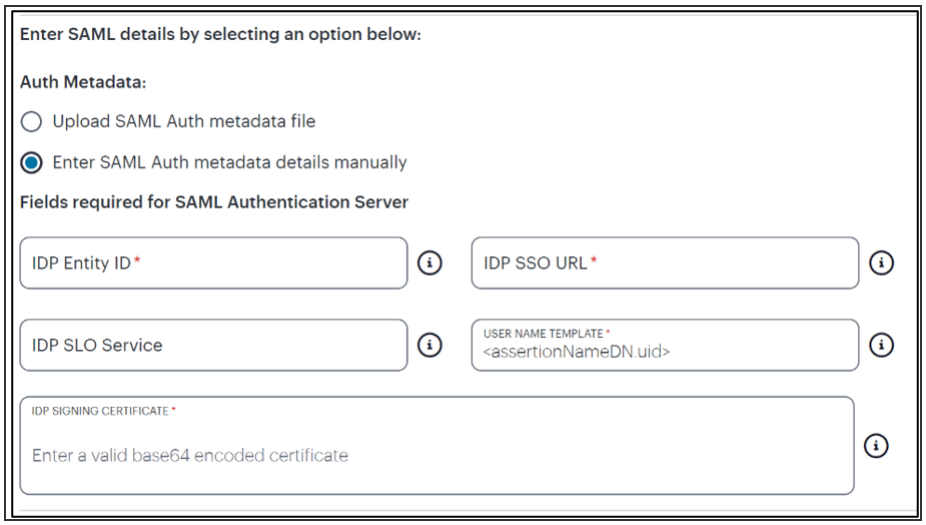

  Configuring SAML (Azure AD) IdP settings manually

  The following minimum settings are required for your SAML authentication service to function correctly. Each setting relates to a value configured in the SAML application on your IdP. Contact your IdP administrator to obtain the relevant details:

  - **IDP Entity ID**: The entity ID is sent as the Issuer value in the SAML assertion generated by the IdP. Specify the Issuer value in assertions generated by the SAML identity provider.
  - **IDP SSO URL**: A URL provisioned by the SAML IdP to support service-provider-initiated Single Sign-On. Use the format `https://<FQDN>`.
  - **IDP SLO Service**: (Optional) A URL to specify the Single Log-Out/sign out endpoint if you want to force re-authentication for increased security. Use the format `https://<FQDN>`. For more information, see Using SAML Single Logout to Force User Authentication.

- **User Name Template**: Specify how the system is to derive the username from the SAML assertion. The default value can be used, or replaced with an alternative specifier that the *Controller* uses from the incoming assertion. For example: `<assertionNameDN.uid>`, the NameID value where *ICS* is the IdP, the UID from X509SubjectName, `<userAttr.attr>`, attr from AttributeStatement attributes.
- **IDP Signing Certificate**: The signing certificate to be used with the SAML app on the IdP. Type or paste in the contents of your Base-64 encoded public key.

- When editing an existing SAML Auth server, the 'Enable Enrollment' option can be enabled or disabled if the SAML Auth server is not being used in any 'User Policy'. If the SAML Auth server is being used in a 'User Policy', then Enable Enrollment button will be grayed out.
- If 'Enable Enrollment' is not selected, then while creation of 'User Policy' of type 'User', the server you have created (without Enable Enrollment) will not be listed.
- When Enrollment is disabled, the enrollment SAML configuration will be deleted. To enable enrollment, you have to again provide enroll SAML auth server configuration.

6. Confirm that your settings are correct, then select **Create Authentication Server** to create the authentication method.

   The new SAML authentication method is added to the list of methods and the process is complete.

7. (Optional) To edit a listed authentication method, click the adjacent three dots, then select **Edit**. Make any required updates and confirm.

8. (Optional) To delete one (or more) an *unused* authentication methods, select the check box for each, then select **Delete**. You must confirm the deletion.

You can now proceed to update the required authentication policy, see Defining Authentication Policies for On-Premises SAML Authentication.

Defining Authentication Policies for On-Premises SAML Authentication

After you have created your SAML authentication method, create or update your authentication policies with the new authentication method.

From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > User Policies**.

The **User Policies** page appears. This page lists all existing user authentication policies.

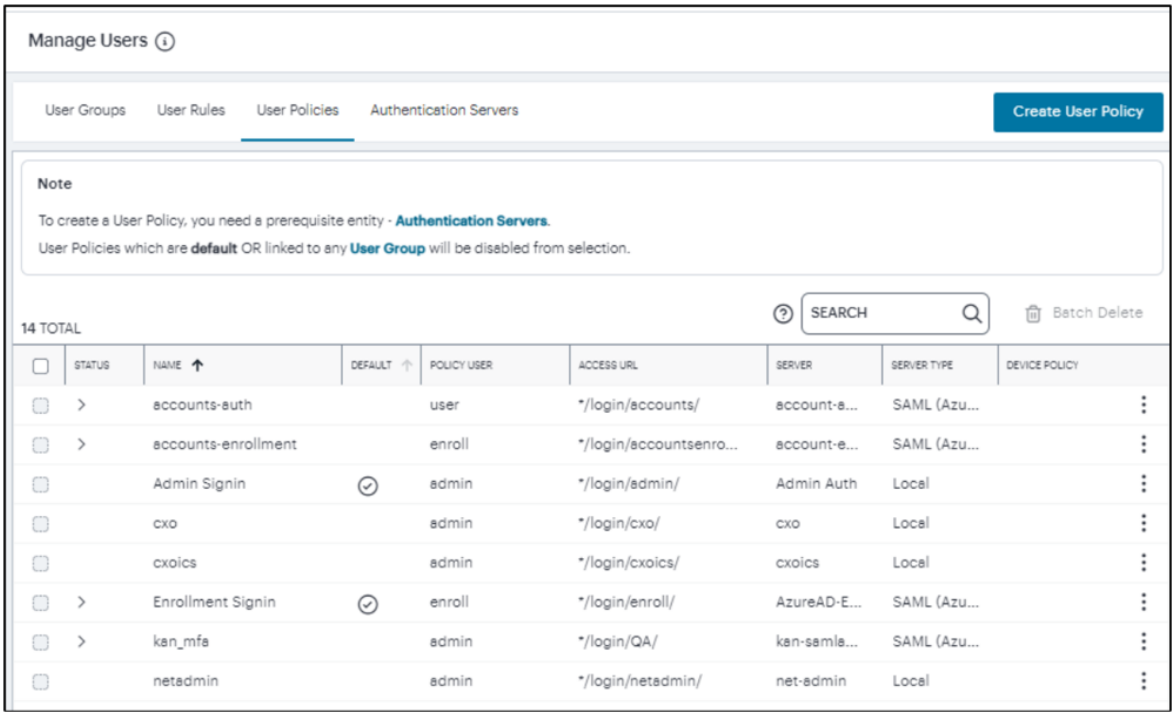

User Authentication Policies

To learn more about the policies on this page, see Viewing User Authentication Policies.

From this page, either create a new custom policy or edit an existing policy. To add a new custom policy:

1. Select **Create User Policy**.

   The **Create Authentication Policy** form appears.

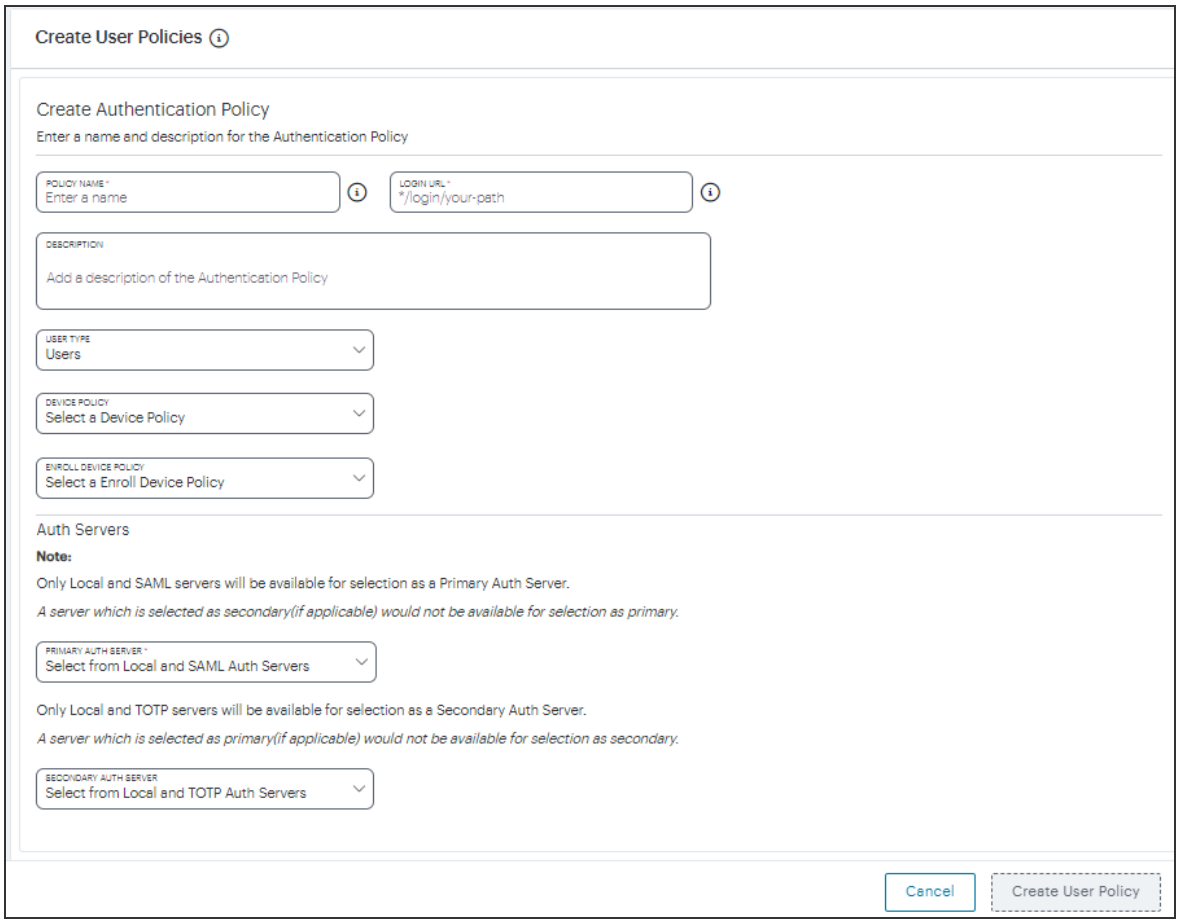

   Add User Authentication

   At any point during this process, you can reset the form data by selecting **Reset Fields**.

   To learn more about how custom policies are used for user login and enrollment, see Adding Custom Authentication Policies.

2. Enter a **Policy Name**.

3. Enter a **Login URL** using the format **/login/<path>/*.

   The URL must start with "**/login/*" and cannot contain any special characters. <path> should be set to a unique value reflecting the endpoint URL you want to define for this authentication policy (appended with a backslash):

   - In the case of user sign-in policies, this is the URL endpoint (appended to the tenant FQDN) to which new users are invited to connect to enroll or sign-in a device with the *Controller*. Example value: "**/login/saleslogin/*".
   - In the case of user enrollment policies, this endpoint identifies the enrollment URL to which users are redirected if they attempt to connect to the equivalent sign-in policy with an un-enrolled device. In most cases, you do not advertise this endpoint to your users. Example value: "**/login/salesenroll/*".

   

   In some enrollment circumstances, such as when using a device pre-installed with an older version of *Ivanti Secure Access Client*, you connect directly to the enrollment policy endpoint to enroll the device. For more details, see Using User Authentication Policies.

4. (Optional) Enter a description for the authentication policy.
5. Select a **User Type** based on the intended authentication activity for this policy. Choose from:

   - **Users**: Select this option to define the user sign-in endpoint for enrolled devices. This is the endpoint that you provide to your users to access the service (regardless of enrollment status).
   - **Administrators**: Select this option to define the authentication endpoint for administrator-level sign-in. This endpoint is used for administrator login to the *Controller* only.

6. (for policies with a **User Type** of "Users" only): Select an **Enroll Device Policy** from the drop-down list to be linked to this sign-in policy (as indicated):

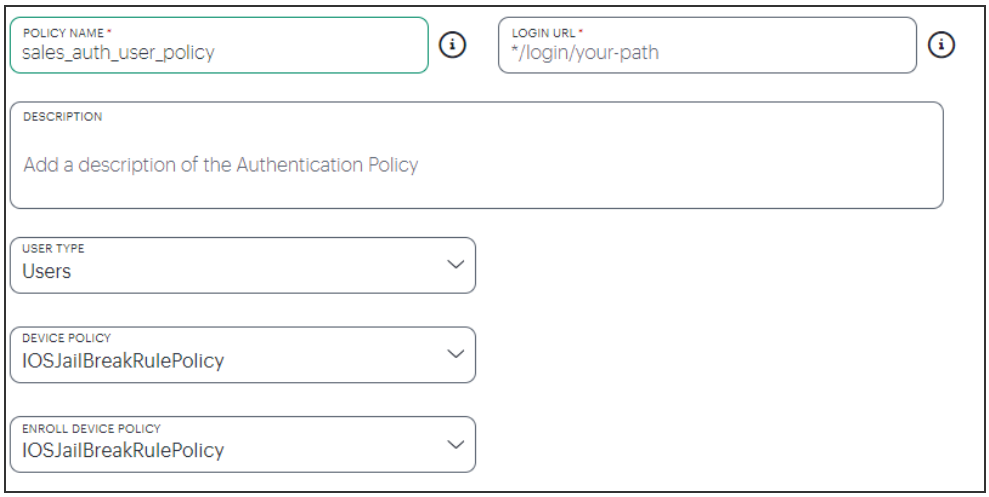

Linking an enrollment policy to a user sign-in policy

This is the enrollment policy to which a user is redirected if it is determined that the device is not yet enrolled. To learn more, see Using User Authentication Policies.

7. Under **Policy Server Details**, select **Primary Auth Server** and choose the required authentication method from the drop-down list:

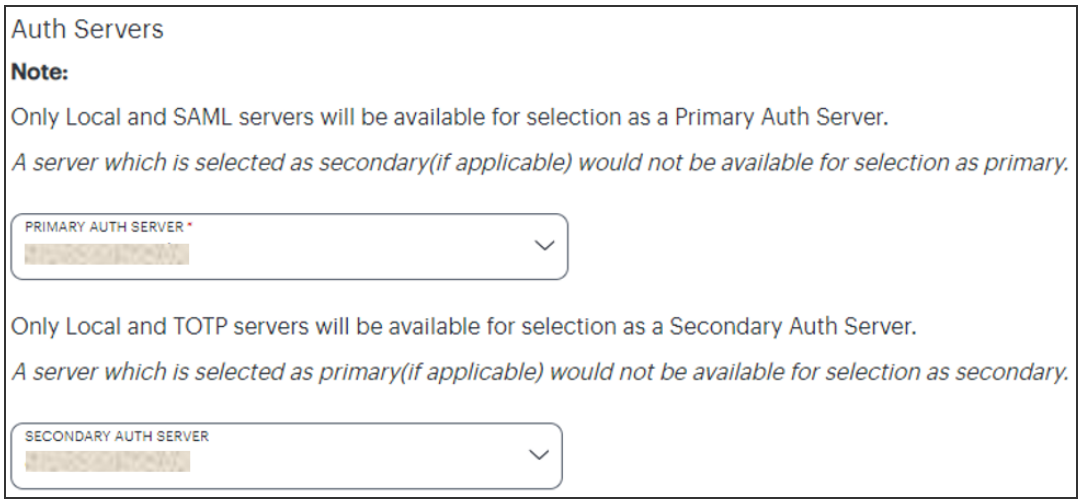

Selecting a primary authentication method for this policy

Alternatively, select *Create Authentication Server* and create a new authentication method as per the steps described earlier in this section.

8. (Optional) Where a secondary method is required for Multi-Factor Authentication, repeat the previous step for **Secondary Auth Server**.

   Secondary authentication methods do not apply to *Enrollment User* type policies. The relevant field is hidden in this case.

9. Select **Add** to create the new policy.

   The new policy is added to the list of authentication policies.

If you instead elect to update an existing custom or built-in policy:

1. Click the three dots adjacent to the relevant policy, then select **Edit**.

   The **Edit authentication policy** form appears.

   For built-in authentication policies, all properties except **Primary Auth Server** and **Secondary Auth Server** (where applicable) are read-only.

2. Set the **Primary Auth Server** to be the new SAML user authentication method (indicated):

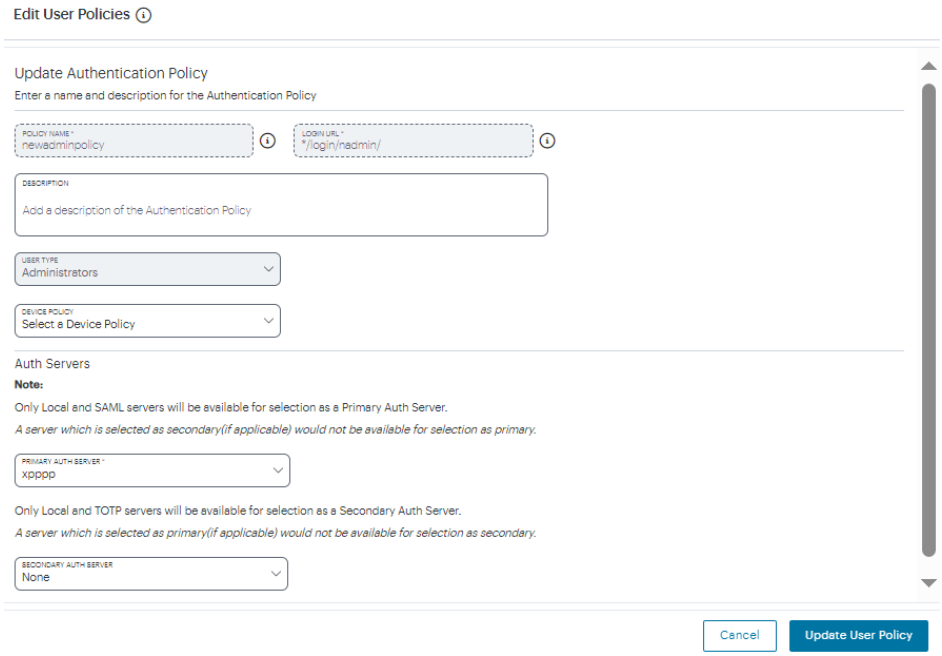

Editing the primary auth server

> SAML authentication can be used only as a Primary Auth Server. If you are using MFA, specify either a *local authentication* or *TOTP* method as the **Secondary Auth Server**.

> If you configure a secondary authentication method in a policy that is currently in use, any active user sessions must be disconnected and reconnected for the changes to take effect.

3. Select **Update User Policy**.

   The list of authentication policies updates.

At this point, the *Controller* uses the uploaded metadata to contact the SAML service. After this process completes, a **Download** function becomes available for the policy. This metadata file is required to configure trusted communication with the remote SAML service. Perform the following steps:

1. Refresh your browser until the **Download** action is visible for the relevant policy.
2. Select the check box for the policy and clear all other check boxes.
3. Select **Download** and save the metadata file.

Repeat these procedures for each required SAML app on your On-Prem *ICS* server. That is, you require separate XML metadata files for your enrollment authentication policy and your sign-in authentication policy.

After you have configured a user authentication policy, you can configure your *ICS* SAML app with the SP Metadata configuration of the *Controller*, see Configuring ICS with Controller Metadata.

Configuring *ICS* with *Controller* Metadata

Before the *Controller* can use a SAML server on an on-premises *ICS* server, you must enable communication between each separate SAML app on the *ICS* server and the *Controller*.

To do this, you must configure the SAML apps on the on-premises *ICS* server with the XML metadata configuration file for the *Controller*.

There are a minimum of two SAML apps:

- A SAML app for user enrollment. For example, called *Enrollment*.
- A SAML app during user sign-in. For example, called *Signin*.

You download the *Controller* XML configuration file when you selected an authentication policy. Alternatively, select the policy, then select **Download** on the **User authentication policies** page, see Viewing User Authentication Policies.

To configure a SAML app on *ICS* with XML from the *Controller*:

1. Log into your *ICS* platform.
2. Navigate to **System > Configuration > SAML**.
3. Select **New Metadata Provider**.
4. Enter a **Name** for the metadata provider.
5. Under **Metadata Provider Location Configuration**:
    - For **Location**, select *Local*.
    - For **Upload Metadata File**, select **Browse** and select the metadata that you saved on your computer in the previous process (see above).

6. Under **Metadata Provider Verification Configuration**:
    - Select the **Accept Unsigned Metadata** check box.
7. Under **Metadata Provider Filter Configuration**:
    - For **Roles**, select the Service Provider check box.
8. Select **Save Changes**.
9. Navigate to **Authentication > Signing In > Sign-In SAML > Identity Provider**.
10. In the **Configuration** section, select **Add SP**.

    The New Peer Service Provider page appears.
11. In the **Service Provider Configuration** and **Certificate Status Checking Configuration** sections, make the necessary service provider specific settings. For more details, refer to the "Configuring Sign-in SAML Identity Provider Settings" section in the "*Ivanti Connect Secure* Administration Guide".
12. In the **Customize IdP Behavior** section, select the **Override Default Configuration** check box.
13. Clear the **Reuse Existing NC (Pulse) Session** check box.
14. Select the **Accept unsigned AuthnRequest** check box.
15. At the bottom of the page, select **Save Changes**.

After the on-premises *ICS* SAML service is configured to connect to the *Controller*, you can then create a user group to apply the newly created authentication policy to your secure applications, see Creating User Rules and User Groups.

Workflow: Creating a SAML Authentication Policy for Okta

nZTA supports the use of Okta as a SAML service to provide authentication for your users.

If you choose to use Okta as a SAML Identity Provider (IdP), you do not create any users locally on the *Controller*. All users will already be present in your remote SAML service.

The process to configure Okta as a custom SAML authentication method within *nZTA* involves the following steps:

1. Create your Okta application and obtain the SAML IdP metadata, see Creating an Okta SAML Application.
2. Define an Okta SAML authentication method in *nZTA* and associate it with your authentication policies, see Defining and Applying Okta Authentication in nZTA.

Configuring *nZTA* to use Okta SAML authentication requires configuration of two separate authentication policies on the *Controller*, *user enrollment* and *user sign-in*. The *Controller* includes built-in default authentication policies for each of these purposes, and also includes the ability to create your own custom policies should you require this authentication mechanism to apply only to a sub-set of your users. Therefore, complete the above steps for each of these policies in turn. Begin with enrollment, and then repeat the process for user sign-in.

To ensure that your users can access the authentication mechanism defined in the policies you configure through this process, make sure your *Secure Access Policies* are configured with a **User Group** in which these authentication policies are defined. To learn more, see Creating/Editing Secure Access Policies.

Before you start, make sure you know the following information:

- The default Assertion Consumer Service (ACS) URL:
    - For the user enrollment policy, the ACS URL uses the form: `https://<tenant FQDN>/dana-na/auth/saml-consumer.cgi`

      (For example, `https://zta1.example.com/dana-na/auth/saml-consumer.cgi`)
    - For the user sign-in policy, the ACS URL uses the form: `https://<tenant MTLS FQDN>/dana-na/auth/saml-consumer.cgi`

      (For example, `https://zta1.e.example.com/dana-na/auth/saml-consumer.cgi`)

- The SP Entity ID of *nZTA* for your integration:
  - For the user enrollment policy, the SP Entity ID uses the form: `https://<tenant FQDN>/dana-na/auth/saml-endpoint.cgi?p=sp1`

    (For example, `https://zta1.example.com/dana-na/auth/saml-endpoint.cgi?p=sp1`)
  - For the user sign-in policy, the SP Entity ID uses the form: `https://<tenant MTLS FQDN>/dana-na/auth/saml-endpoint.cgi?p=sp2`

    (For example, `https://zta1.e.example.com/dana-na/auth/saml-endpoint.cgi?p=sp2`)

Creating an Okta SAML Application

- To fully configure Okta as a SAML authenticator in *nZTA*, complete these steps for each of your user enrollment and user sign-in policies in turn.
- For the latest configuration details, see Okta Documentation.

To create a new Okta application and obtain the SAML IdP Metadata file:

1. Log in to your Okta account and select *Classic UI* mode.
2. In the Admin Console, navigate to **Applications** > **Applications**.
3. Select **Add Application**.
4. Start the *Application Integration Wizard* by selecting **Create New App**.
5. For **Platform**, select "Web".
6. For **Sign-on method**, select "SAML 2.0".
7. Select **Create**.
8. In the *General Settings* tab, enter an application name.

   

   Ivanti advises using a descriptive name that relates to the user authentication policy this application is created for.
9. (Optional) Upload a logo for the application.

10. In the *Configure SAML* tab, enter the following information corresponding to the authentication policy for which you are creating this application.

    - **Single sign on URL**: Enter your ACS URL
    - **Audience URI (SP Entity ID)**: Enter your *nZTA* Entity ID
    - **Name ID Format**: Select "Email Address"
    - **Application username**: Select "Okta username".
    - (Optional) If you want to use Okta group membership policies to categorize users, complete the **Fill in Group Attribute Statements** section to filter users by group membership in the SAML assertion.

      For example, to set a filter that returns all groups to which the user is assigned, you might enter the following details:

      - **Name**: ZTAGroups
      - **Filter**: Matches Regex
      - **Value**: .*

      For more information, see Using Okta Group Attribute Statements to Categorize Users.
    - For all other fields, use the default values.

11. Select **Next** to continue.
12. For **Are you a customer or partner?**, select the option that is most applicable.
13. Select **Finish**.

    The *Settings* page for your application appears.

14. In the *Sign On* tab, download the **Identity Provider metadata** file. Save this file to your local workstation.

15. In the *Applications* tab, select **Assign Applications**. Then select **Assign to People** or **Assign to Groups** as per your requirements.
    - For each user or group you want to assign to your application, select **Assign** against the respective item.
    - After you have completed your assignments, select **Done**.
16. Proceed to define a new authentication method in *nZTA*, see Defining and Applying Okta Authentication in nZTA.

Using Okta Group Attribute Statements to Categorize Users

> This section is applicable only to scenarios where Okta SAML authentication is augmented with user categorization through group attribute policies.

To configure Okta with group membership policies, and to enable *nZTA* to query these policies when authenticating access requests, add *Group Attribute Statements* to your Okta user enrollment and user sign-in applications. Then, update your enrollment and sign-in *User Rules* within *nZTA* to use the corresponding *Custom* SAML attributes.

The following table lists common Okta *Group Attribute Statements*, designed to return matching user groups in a SAML assertion:

> The attribute name "ZTAGroups" is used here as an example. Use the attribute tag which corresponds to your Okta group configuration.

Okta Group Attribute Name	Okta Group Attribute Filter	Groups Returned
ZTAGroups	Contains: ZTA	All groups containing "ZTA" (ignores case)
ZTAGroups	StartsWith: ZTA	All groups starting with "ZTA" (ignores case)
ZTAGroups	Equals: ZTAUsers	The "ZTAUsers" group only
ZTAGroups	Matches regex: .*	All groups to which the user is assigned

The following are examples of Okta custom SAML expressions for use within *nZTAUser Rules*:

- Users that are members of only the named group:

```
userAttr.ZTAGroups = "ZTAGroup1"
```

- Users that are members of either named group:

```
userAttr.ZTAGroups = ("ZTAGroup1" or "ZTAGroup2")
```

- Users that are members of either named group, or both groups:

```
samlMultiValAttr.ZTAGroups = ("ZTAGroup1" or "ZTAGroup2")
```

- Users that are only members of both named groups:

```
samlMultiValAttr.ZTAGroups = ("ZTAGroup1" and "ZTAGroup2")
```

- Users that are members of "ZTAGroup1", but not members of "ZTAGroup2":

```
 samlMultiValAttr.ZTAGroups = "ZTAGroup1" and
samlMultiValAttr.ZTAGroups != "ZTAGroup2"
```

- Users that are members of either "ZTAGroup1" or "ZTAGroup2", or both groups, and also a member of "ZTAGroup3":

```
samlMultiValAttr.ZTAGroups = ("ZTAGroup1" or "ZTAGroup2") and
samlMultiValAttr.ZTAGroups = "ZTAGroup3"
```

To learn more about adding custom SAML user rules, see Creating User Rules.

Defining and Applying Okta Authentication in *nZTA*

To fully configure Okta as a SAML authenticator in *nZTA*, complete these steps for each of your user enrollment and user sign-in policies in turn.

To define a new authentication method using Okta as the SAML IdP:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.

   The **Network Overview** page appears.

2. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > Authentication Server**.

   The *Authentication Server* page appears. This page lists all existing user authentication methods.

3. To add a new custom SAML authentication method, select **Create Authentication Server**.

   The **Create Authentication Server** form appears:

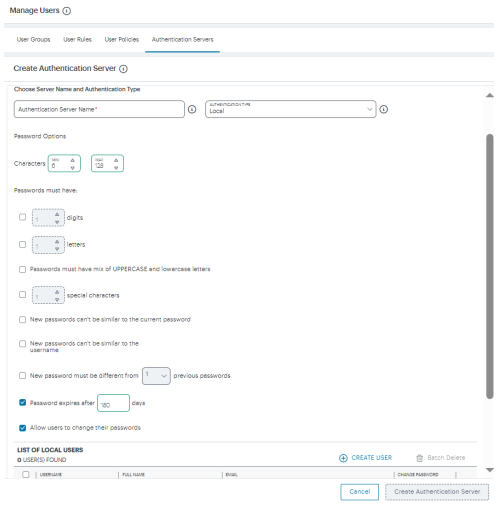

   Creating an authentication server

> At any point during this process, you can reset the form data by selecting **Reset Fields**.

4. Under **Choose Server Name and Authentication Type**:

   - Select the **Authentication Type** of *SAML (Custom)*.

     The form expands to show additional settings:

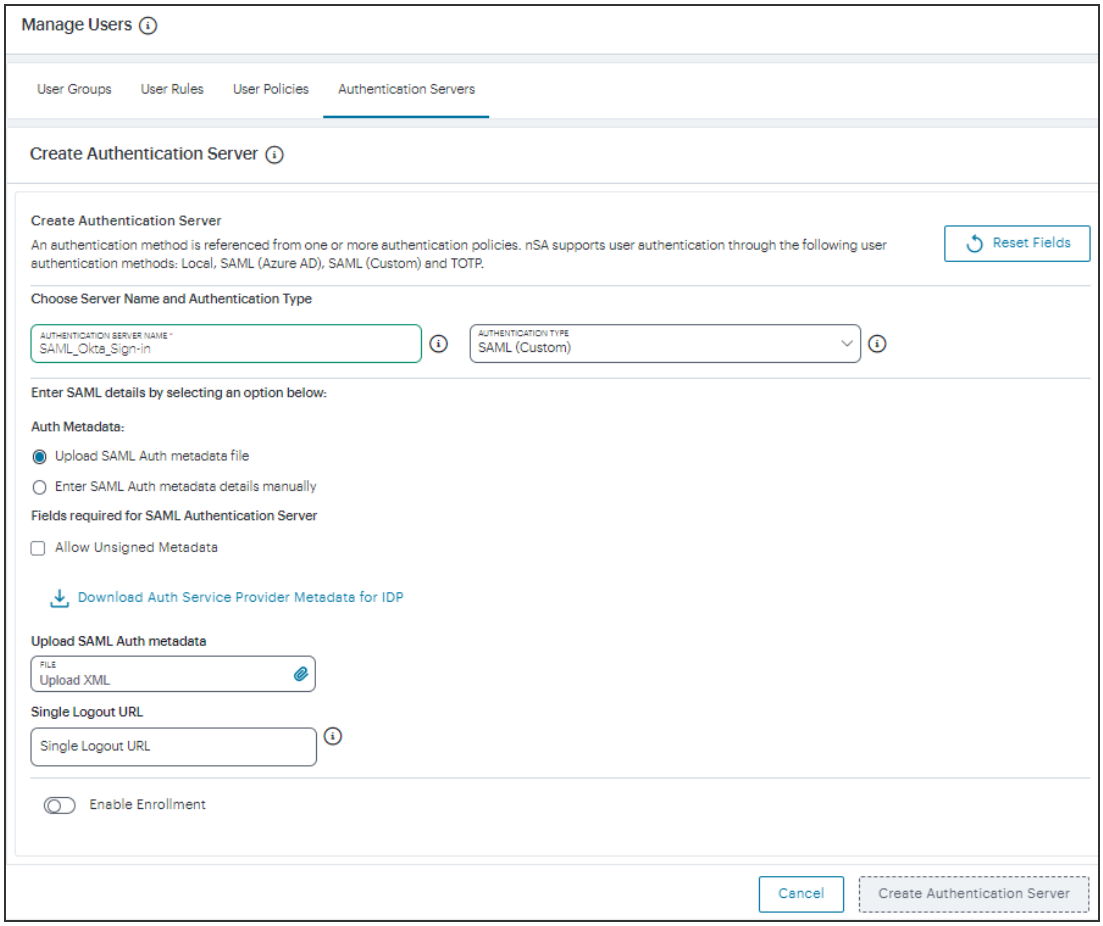

     Configuring Okta SAML authentication settings

   - Specify an **Authentication Server Name**. For example: *SAML_Okta_Enroll* or *SAML_Okta_SignIn*.

5. (Optional) Specify a **Single Logout URL**. For more information, see Using SAML Single Logout to Force User Authentication.

6. To provide your SAML IdP settings, select one of the following:

   - Select **Upload SAML Auth metadata file** to upload a digitally-signed (or unsigned) Okta Identity Provider (IdP) metadata file, as obtained while creating a new Okta SAML application for this activity. That is, for either user enrollment or user sign-in. See Creating an Okta SAML Application.

     The **Download Auth Service Provider Metadata for IDP** link is enabled.

     1. (Optional) Specify a **Single Logout URL**. For more information, see Using SAML Single Logout to Force User Authentication.

        

        By default, the *Controller* expects a signed metadata definition file. To allow an unsigned metadata file, select **Allow Unsigned Metadata**.

     2. Click the **Download Auth Service Provider Metadata for IDP** link. Retain the downloaded file for later use.

- Select **Enter SAML Auth metada details manually** to manually enter the required IdP SAML settings. Use this option in scenarios where a SAML federation metadata file is not available or incomplete.

  Then, enter the following details:

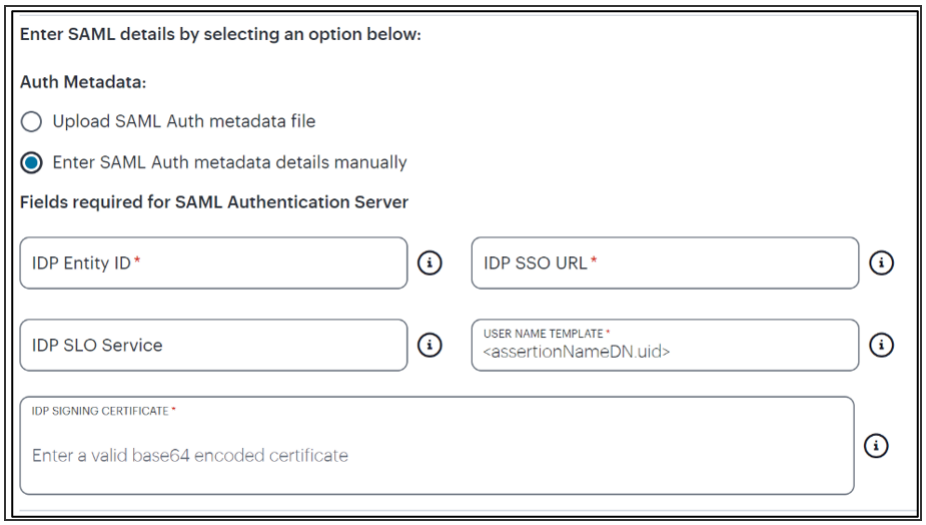

  Configuring SAML (Custom) IdP settings manually

  The following minimum settings are required for your SAML authentication service to function correctly. Each setting relates to a value configured in the SAML application on your IdP. Contact your IdP administrator to obtain the relevant details:

  - **IDP Entity ID**: The entity ID is sent as the Issuer value in the SAML assertion generated by the IdP. Specify the Issuer value in assertions generated by the SAML identity provider.
  - **IDP SSO URL**: A URL provisioned by the SAML IdP to support service-provider-initiated Single Sign-On. Use the format `https://<FQDN>`.
  - **IDP SLO Service**: (Optional) A URL to specify the Single Log-Out/sign out endpoint if you want to force re-authentication for increased security. Use the format `https://<FQDN>`. For more information, see Using SAML Single Logout to Force User Authentication.

- **User Name Template**: Specify how the system is to derive the username from the SAML assertion. The default value can be used, or replaced with an alternative specifier that the *Controller* uses from the incoming assertion. For example: `<assertionNameDN.uid>`, the NameID value where *ICS* is the IdP, the UID from X509SubjectName, `<userAttr.attr>`, attr from AttributeStatement attributes.
- **IDP Signing Certificate**: The signing certificate to be used with the SAML app on the IdP. Type or paste in the contents of your Base-64 encoded public key.

If, at a later date, you need to replace the metadata definition file with a modified version, edit the authentication method through the *User Authentication* page and repeat this step. Either edit the existing metadata file and re-upload, or replace it completely with a new version. In both cases, however, make sure your metadata file is valid before uploading it through this process.

7. If this Auth server is used with User Policy of type "User", then click **Enable Enrollment**.

Enable Enrollment

8. To provide your SAML IdP settings, select one of the following:

    - Select **Upload SAML Enroll metadata file** if not selected already. This is selected by default.

        1. (Optional) Specify a **Single Logout URL**. For more information, see Using SAML Single Logout to Force User Authentication.

            

            By default, the *Controller* expects a signed metadata definition file. To allow an unsigned metadata file, select **Allow Unsigned Metadata**.

        2. Click the **Download Enroll Service Provider Metadata for IDP** link.

- Select **Enter SAML Enroll metada details manually** to manually enter the required IdP SAML settings. Use this option in scenarios where a SAML federation metadata file is not available or incomplete.

  Then, enter the following details:

  

  Configuring SAML (Azure AD) IdP settings manually

  The following minimum settings are required for your SAML authentication service to function correctly. Each setting relates to a value configured in the SAML application on your IdP. Contact your IdP administrator to obtain the relevant details:

  - **IDP Entity ID**: The entity ID is sent as the Issuer value in the SAML assertion generated by the IdP. Specify the Issuer value in assertions generated by the SAML identity provider.
  - **IDP SSO URL**: A URL provisioned by the SAML IdP to support service-provider-initiated Single Sign-On. Use the format `https://<FQDN>`.
  - **IDP SLO Service**: (Optional) A URL to specify the Single Log-Out/sign out endpoint if you want to force re-authentication for increased security. Use the format `https://<FQDN>`. For more information, see Using SAML Single Logout to Force User Authentication.
  - **User Name Template**: Specify how the system is to derive the username from the SAML assertion. The default value can be used, or replaced with an alternative specifier that the *Controller* uses from the incoming assertion. For example: `<assertionNameDN.uid>`, the NameID value where *ICS* is the IdP, the UID from X509SubjectName, `<userAttr.attr>`, attr from AttributeStatement attributes.
  - **IDP Signing Certificate**: The signing certificate to be used with the SAML app on the IdP. Type or paste in the contents of your Base-64 encoded public key.

- When editing an existing SAML Auth server, the 'Enable Enrollment' option can be enabled or disabled if the SAML Auth server is not being used in any 'User Policy'. If the SAML Auth server is being used in a 'User Policy', then Enable Enrollment button will be grayed out.
- If 'Enable Enrollment' is not selected, then while creation of 'User Policy' of type 'User', the server you have created (without Enable Enrollment) will not be listed.
- When Enrollment is disabled, the enrollment SAML configuration will be deleted. To enable enrollment, you have to again provide enroll SAML auth server configuration.

9. Confirm that your settings are correct, then select **Create Authentication Server** to create the authentication method.

   The *Authentication Server* page lists the new Okta authentication method.

After you have created your Okta authentication method, create or update your authentication policies with the new authentication method:

1. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > User Policies**.

   The **User Policies** page appears. This page lists all existing user authentication policies.

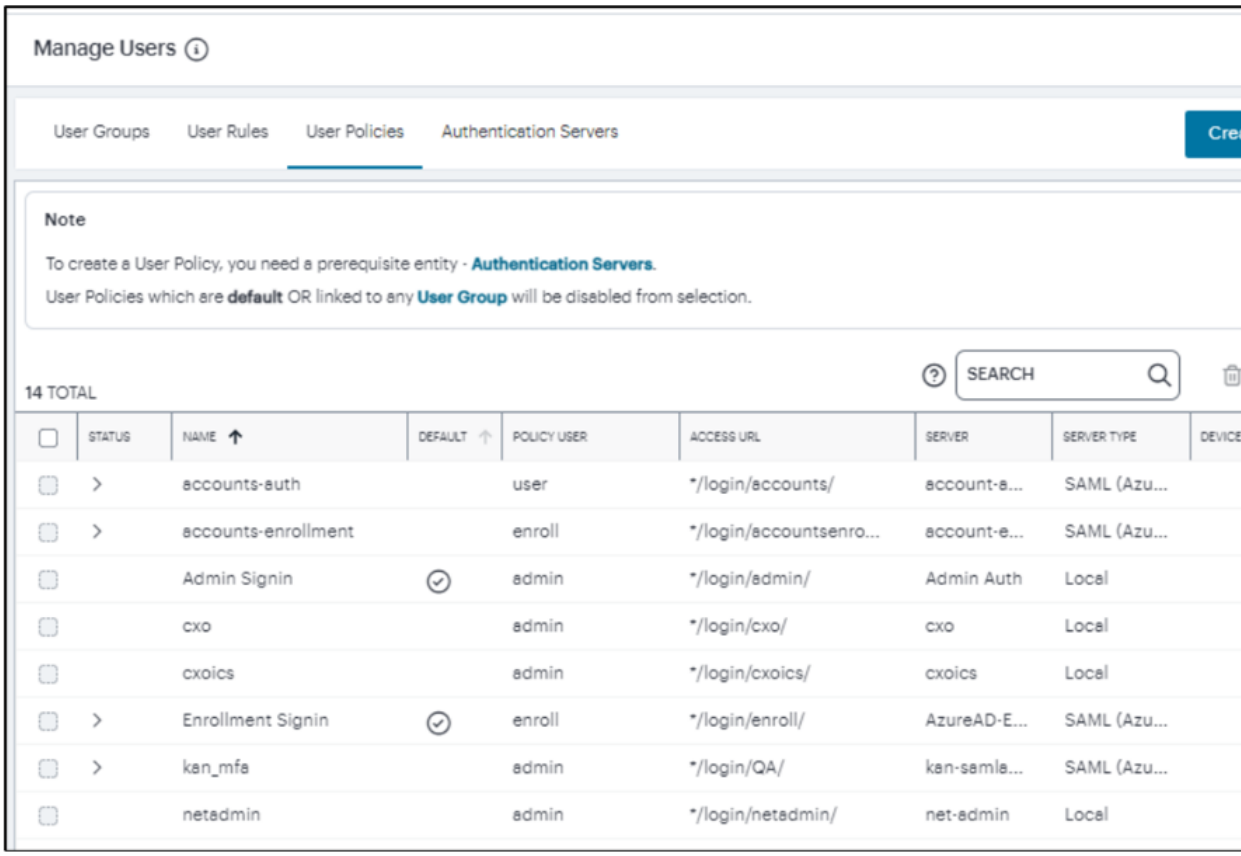

   User Authentication Policies

   To learn more about the policies on this page, see Viewing User Authentication Policies.

   From this page, either create a new custom policy or edit an existing policy.

2. To add a new custom policy, select **Create User Policy**.

   The **Create Authentication Policy** form appears.

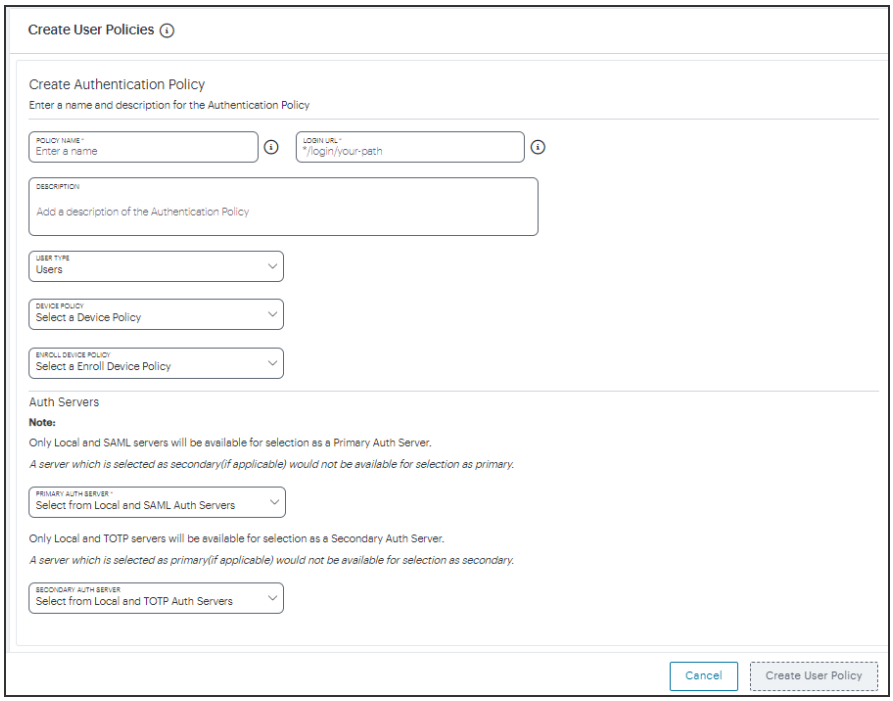

   Add User Authentication

> At any point during this process, you can reset the form data by selecting **Reset Fields**.

> To learn more about how custom policies are used for user login and enrollment, see Adding Custom Authentication Policies.

3. Enter a **Policy Name**.

4. Enter a **Login URL** using the format **/login/<path>/*.

   The URL must start with "**/login/*" and cannot contain any special characters. <path> should be set to a unique value reflecting the endpoint URL you want to define for this authentication policy (appended with a backslash):

   - In the case of user sign-in policies, this is the URL endpoint (appended to the tenant FQDN) to which new users are invited to connect to enroll or sign-in a device with the *Controller*. Example value: "**/login/saleslogin/*".
   - In the case of user enrollment policies, this endpoint identifies the enrollment URL to which users are redirected if they attempt to connect to the equivalent sign-in policy with an un-enrolled device. In most cases, you do not advertise this endpoint to your users. Example value: "**/login/salesenroll/*".

   

   In some enrollment circumstances, such as when using a device pre-installed with an older version of *Ivanti Secure Access Client*, you connect directly to the enrollment policy endpoint to enroll the device. For more details, see Using User Authentication Policies.

5. (Optional) Enter a description for the authentication policy.
6. Select a **User Type** based on the intended authentication activity for this policy. Choose from:

   - **Users**: Select this option to define the user sign-in endpoint for enrolled devices. This is the endpoint that you provide to your users to access the service (regardless of enrollment status).
   - **Administrators**: Select this option to define the authentication endpoint for administrator-level sign-in. This endpoint is used for administrator login to the *Controller* only.

7. (for policies with a **User Type** of "Users" only): Select an **Enroll Device Policy** from the drop-down list to be linked to this sign-in policy (as indicated):

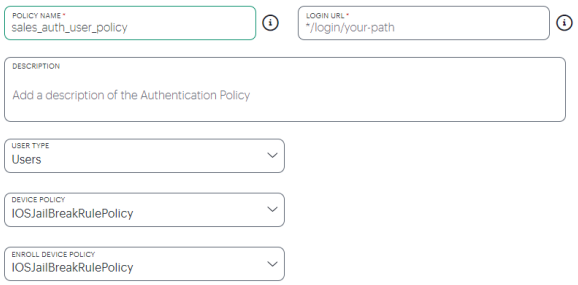

Linking an enrollment policy to a user sign-in policy

This is the enrollment policy to which a user is redirected if it is determined that the device is not yet enrolled. To learn more, see Using User Authentication Policies.

8. Under **Policy Server Details**, select **Primary Auth Server** and choose the required authentication method from the drop-down list:

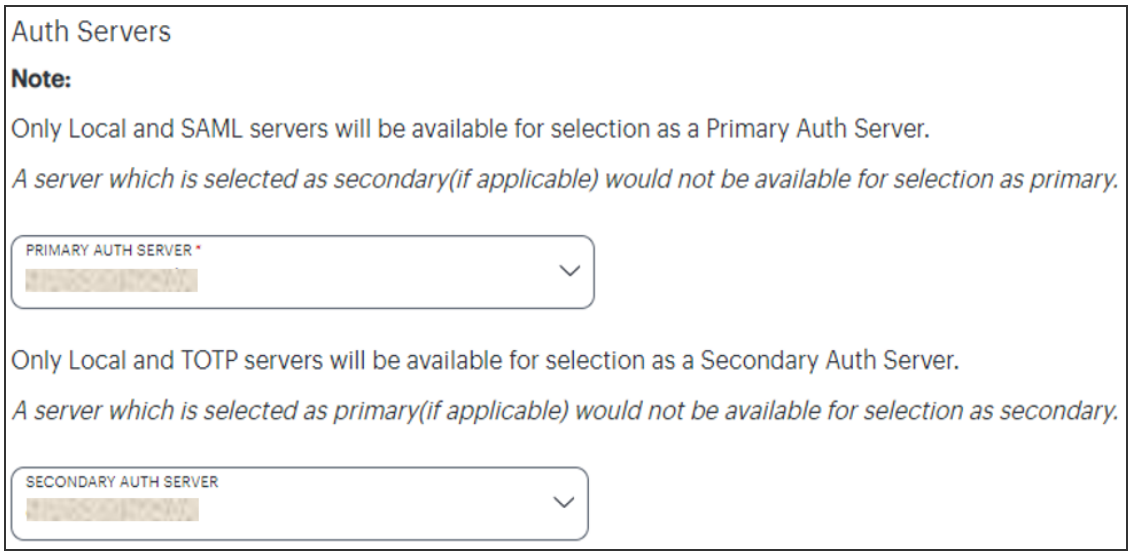

Selecting a primary authentication method for this policy

Alternatively, select *Add New Server* and create a new authentication method as per the steps described earlier in this section.

9. (Optional) Where a secondary method is required for Multi-Factor Authentication, repeat the previous step for **Secondary Auth Server**.

   Secondary authentication methods do not apply to *Enrollment User* type policies. The relevant field is hidden in this case.

10. Select **Add** to create the new policy.

    The new policy is added to the list of authentication policies.

If you instead elect to update an existing custom or built-in policy:

1. Click the three dots adjacent to the relevant policy, then select **Edit**.

   The **Update Authentication Policy** form appears.

   For built-in authentication policies, all properties except **Primary Auth Server** are read-only.

2. Set the **Primary Auth Server** to be the new Okta SAML user authentication method:

   SAML authentication can be used only as a Primary Auth Server. If you are using MFA, specify either a *local authentication* or *TOTP* method as the **Secondary Auth Server**.

   If you configure a secondary authentication method in a policy that is currently in use, any active user sessions must be disconnected and reconnected for the changes to take effect.

3. Select **Update User Policy**.

   The list of authentication policies updates.

Complete these steps for each of your user enrollment and user sign-in policies in turn.

Workflow: Creating a SAML Authentication Policy for Ping Identity

nZTA supports the use of Ping Identity (PingID) as a SAML service to provide authentication for your users.

If you choose to use PingID as a SAML Identity Provider (IdP), you do not create any users locally on the *Controller*. All users will already be present in your remote SAML service.

The process to configure PingID as a custom SAML authentication method within *nZTA* involves the following steps:

1. Create your PingID application and obtain the SAML IdP metadata, see Creating a PingID SAML Application.
2. Define a PingID SAML authentication method in *nZTA* and associate it with your authentication policies, see Defining and Applying PingID Authentication in nZTA.
3. Obtain the *nZTA* Service Provider (SP) metadata and upload it back to the PingID application, see Updating Your PingID SAML Application with Your SP Metadata.

Configuring *nZTA* to use PingID SAML authentication requires configuration of two separate authentication policies on the *Controller*, *user enrollment* and *user sign-in*. The *Controller* includes built-in default authentication policies for each of these purposes, and also includes the ability to create your own custom policies should you require this authentication mechanism to apply only to a sub-set of your users. Therefore, complete the above steps for each of these policies in turn. Begin with enrollment, and then repeat the process for user sign-in.

To ensure that your users can access the authentication mechanism defined in the policies you configure through this process, make sure your *Secure Access Policies* are configured with a **User Group** in which these authentication policies are defined. To learn more, see Creating/Editing Secure Access Policies.

Creating a PingID SAML Application

- To fully configure PingID as a SAML authenticator in *nZTA*, complete these steps for each of your user enrollment and user sign-in policies in turn.
- For the latest configuration details, see PingIdentity Documentation.

To create a new PingID application and obtain the SAML IdP Metadata file:

1. Log in to the *PingOne* web portal (https://admin.pingone.com) and navigate to **My Applications**.

2. Select **Add Application** > **New SAML Application**:

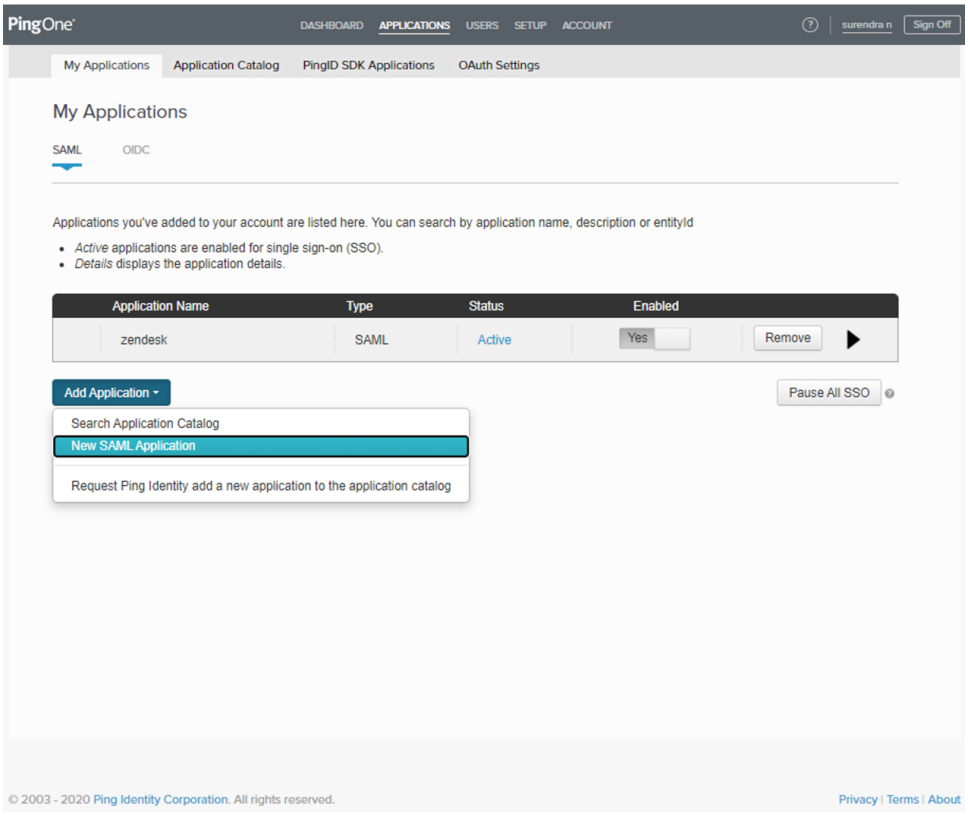

Adding a PingID application

3. In the *Application Details* step, enter an **Application Name** and **Application Description** for the new application:

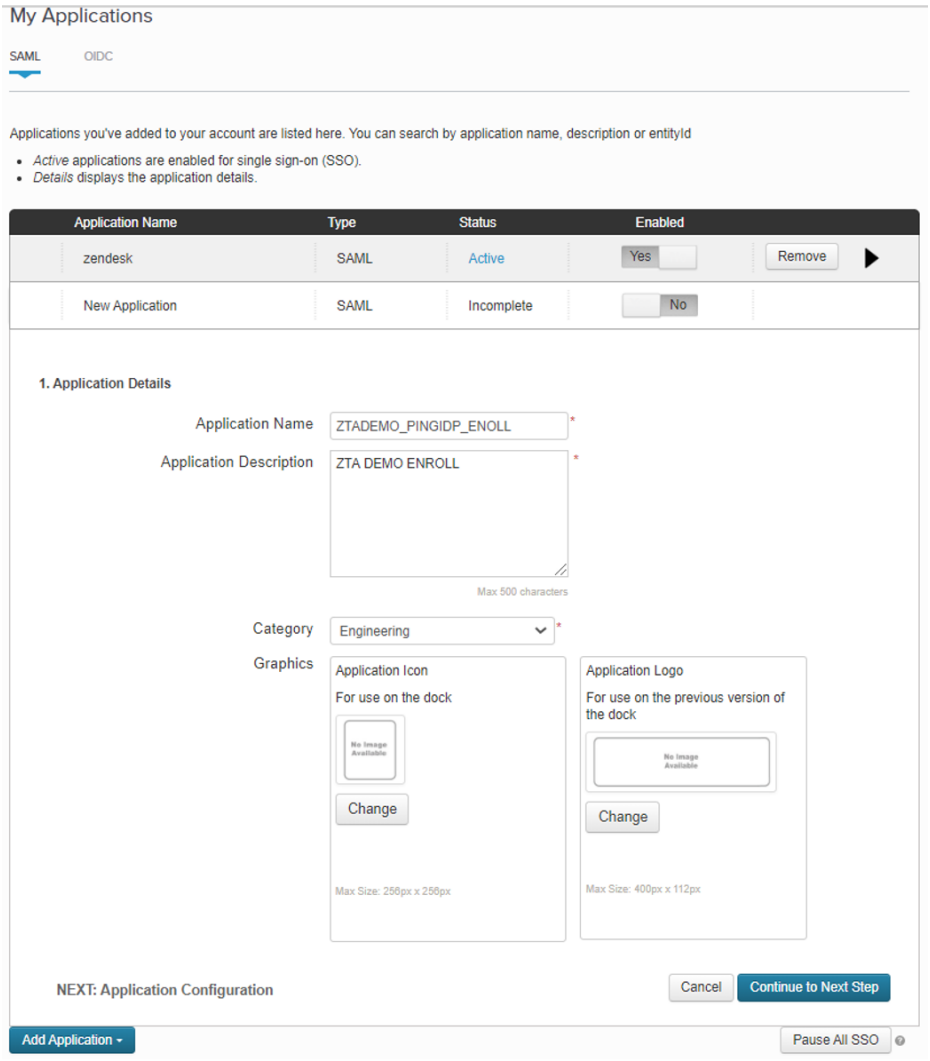

Entering a name and description for your new PingID application

Ivanti advises using a descriptive name that relates to the user authentication policy this application is created for.

4. Select **Continue to Next Step**.

5. In the *Application Configuration* step, locate the **SAML Metadata** field and select the adjacent **Download** link to obtain the SAML metadata file. Save this file to your local workstation:

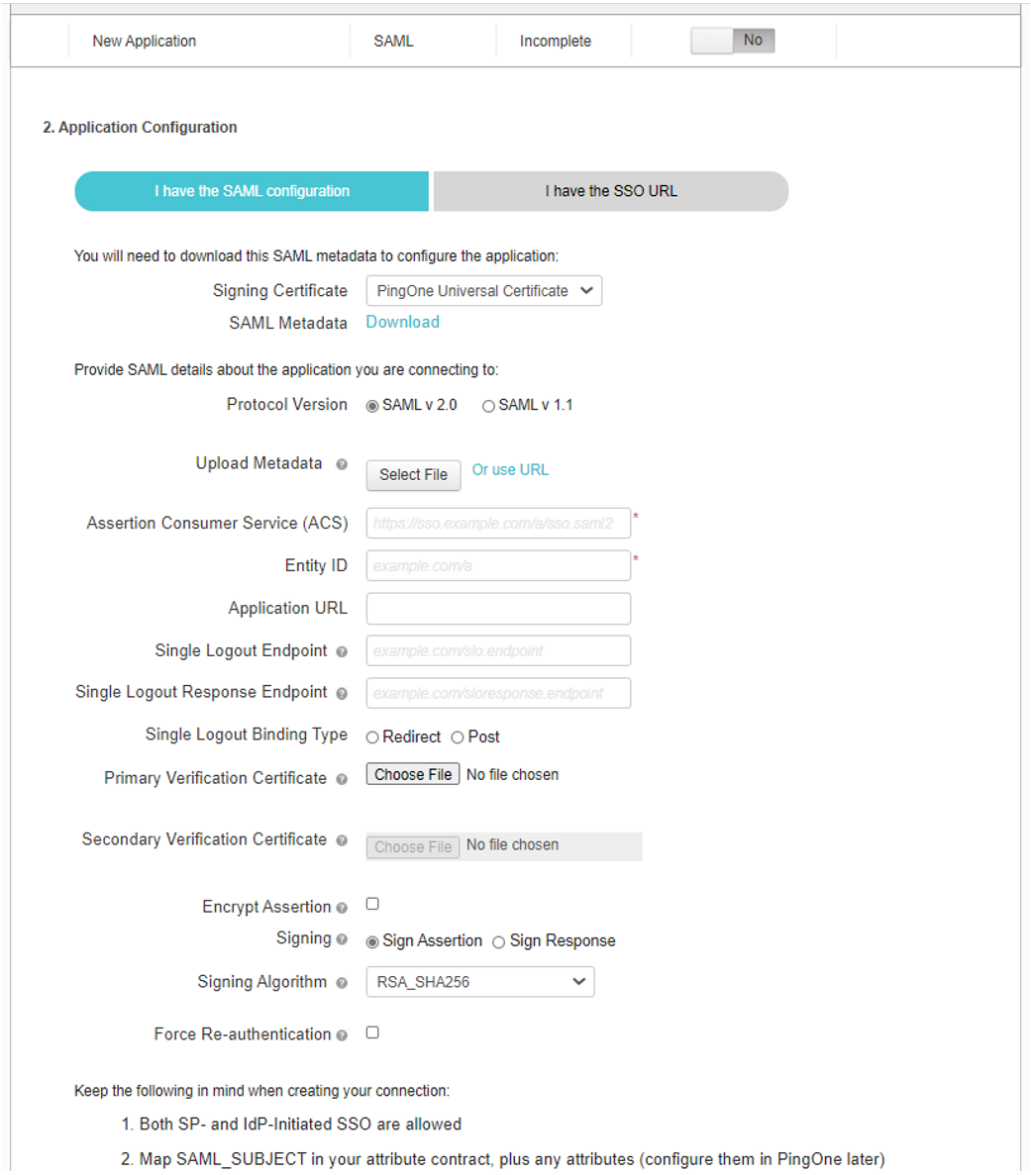

Downloading the PingID application IdP metadata file

6. Keep this browser page open in order to finish creating the application after you have obtained the SP metadata file from *nZTA*.

7. Proceed to define a new authentication method in *nZTA*, see Defining and Applying PingID Authentication in nZTA.

Defining and Applying PingID Authentication in *nZTA*

To fully configure PingID as a SAML authenticator in *nZTA*, complete these steps for each of your user enrollment and user sign-in policies in turn.

To define a new authentication method using PingID as the SAML IdP:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.

   The *Network Overview* page appears.

2. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > Authentication Servers**.

   The *Authentication Servers* page appears. This page lists all existing user authentication methods.

3. To add a new custom SAML authentication method, select **Create Authentication Server**.

   The **Create Authentication Server** form appears:

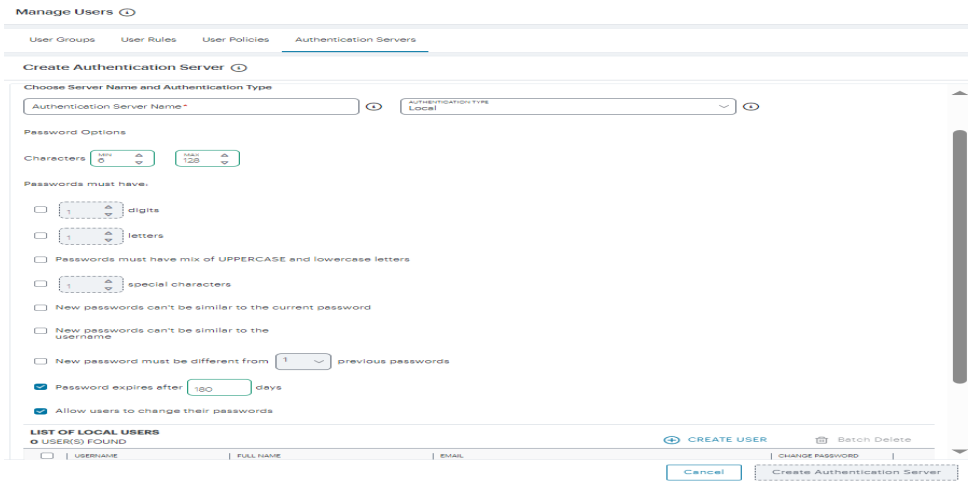

   Adding a user authentication method

   At any point during this process, you can reset the form data by selecting **Reset Fields**.

4. Under **Choose Authentication Name and Type**:

   - Select the **Authentication Type** of *SAML (Custom)*.

     The form expands to show additional settings:

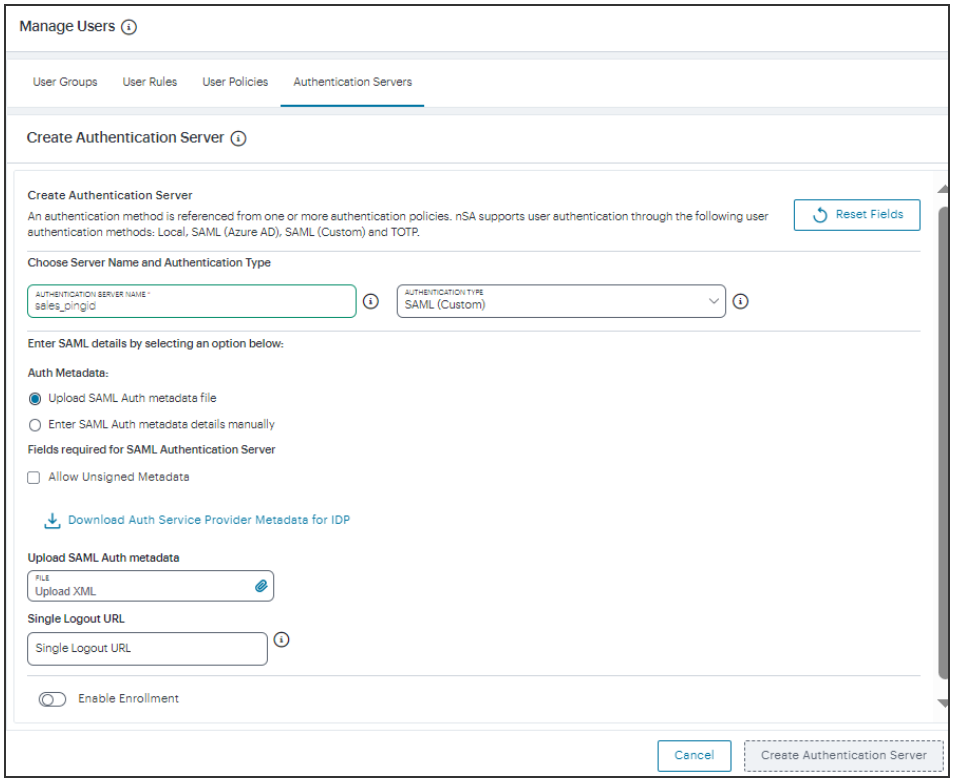

     Configuring PingID SAML authentication settings

   - Specify an **Authentication Server Name**. For example: *SAML_Ping_Enroll* or *SAML_Ping_SignIn*.

5. To provide your SAML IdP settings, select one of the following:

   - Select **Upload SAML Auth metadata file** to upload a digitally-signed (or unsigned) PingID Identity Provider (IdP) metadata file, as obtained while creating a new PingID SAML application for this activity. That is, for either user enrollment or user sign-in. See Creating a PingID SAML Application.

     The **Download Auth Service Provider Metadata for IDP** link is enabled.

     1. (Optional) Specify a **Single Logout URL**. For more information, see Using SAML Single Logout to Force User Authentication.

        

        By default, the *Controller* expects a signed metadata definition file. To allow an unsigned metadata file, select **Allow Unsigned Metadata**.

     2. Click the **Download Auth Service Provider Metadata for IDP** link. Retain the downloaded file for later use.

- Select **Enter SAML Auth metada details manually** to manually enter the required IdP SAML settings. Use this option in scenarios where a SAML federation metadata file is not available or incomplete.

  Then, enter the following details:

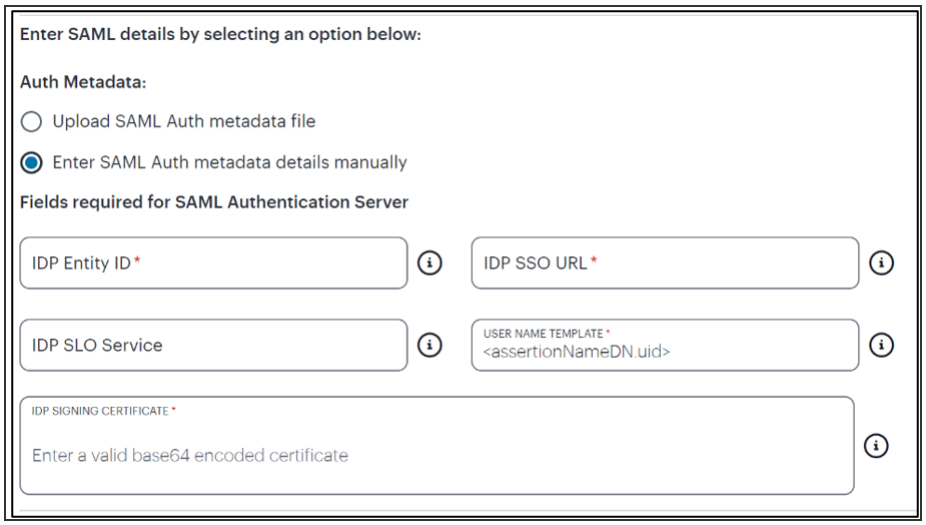

  Configuring SAML (Custom) IdP settings manually

  The following minimum settings are required for your SAML authentication service to function correctly. Each setting relates to a value configured in the SAML application on your IdP. Contact your IdP administrator to obtain the relevant details:

  - **IDP Entity ID**: The entity ID is sent as the Issuer value in the SAML assertion generated by the IdP. Specify the Issuer value in assertions generated by the SAML identity provider.
  - **IDP SSO URL**: A URL provisioned by the SAML IdP to support service-provider-initiated Single Sign-On. Use the format `https://<FQDN>`.
  - **IDP SLO Service**: (Optional) A URL to specify the Single Log-Out/sign out endpoint if you want to force re-authentication for increased security. Use the format `https://<FQDN>`. For more information, see Using SAML Single Logout to Force User Authentication.

- **User Name Template**: Specify how the system is to derive the username from the SAML assertion. The default value can be used, or replaced with an alternative specifier that the *Controller* uses from the incoming assertion. For example: `<assertionNameDN.uid>`, the NameID value where *ICS* is the IdP, the UID from X509SubjectName, `<userAttr.attr>`, attr from AttributeStatement attributes.
- **IDP Signing Certificate**: The signing certificate to be used with the SAML app on the IdP. Type or paste in the contents of your Base-64 encoded public key.

If, at a later date, you need to replace the metadata definition file with a modified version, edit the authentication method through the *User Authentication* page and repeat this step. Either edit the existing metadata file and re-upload, or replace it completely with a new version. In both cases, however, make sure your metadata file is valid before uploading it through this process.

6. If this Auth server is used with User Policy of type “User”, then click **Enable Enrollment**.

Enable Enrollment

7. To provide your SAML IdP settings, select one of the following:

    - Select **Upload SAML Enroll metadata file** if not selected already. This is selected by default.

        1. (Optional) Specify a **Single Logout URL**. For more information, see Using SAML Single Logout to Force User Authentication.

            By default, the *Controller* expects a signed metadata definition file. To allow an unsigned metadata file, select **Allow Unsigned Metadata**.

        2. Click the **Download Enroll Service Provider Metadata for IDP** link.

- Select **Enter SAML Enroll metada details manually** to manually enter the required IdP SAML settings. Use this option in scenarios where a SAML federation metadata file is not available or incomplete.

  Then, enter the following details:

  

  Configuring SAML (Azure AD) IdP settings manually

  The following minimum settings are required for your SAML authentication service to function correctly. Each setting relates to a value configured in the SAML application on your IdP. Contact your IdP administrator to obtain the relevant details:

  - **IDP Entity ID**: The entity ID is sent as the Issuer value in the SAML assertion generated by the IdP. Specify the Issuer value in assertions generated by the SAML identity provider.
  - **IDP SSO URL**: A URL provisioned by the SAML IdP to support service-provider-initiated Single Sign-On. Use the format `https://<FQDN>`.
  - **IDP SLO Service**: (Optional) A URL to specify the Single Log-Out/sign out endpoint if you want to force re-authentication for increased security. Use the format `https://<FQDN>`. For more information, see Using SAML Single Logout to Force User Authentication.
  - **User Name Template**: Specify how the system is to derive the username from the SAML assertion. The default value can be used, or replaced with an alternative specifier that the *Controller* uses from the incoming assertion. For example: `<assertionNameDN.uid>`, the NameID value where *ICS* is the IdP, the UID from X509SubjectName, `<userAttr.attr>`, attr from AttributeStatement attributes.
  - **IDP Signing Certificate**: The signing certificate to be used with the SAML app on the IdP. Type or paste in the contents of your Base-64 encoded public key.

- When editing an existing SAML Auth server, the 'Enable Enrollment' option can be enabled or disabled if the SAML Auth server is not being used in any 'User Policy'. If the SAML Auth server is being used in a 'User Policy', then Enable Enrollment button will be grayed out.
- If 'Enable Enrollment' is not selected, then while creation of 'User Policy' of type 'User', the server you have created (without Enable Enrollment) will not be listed.
- When Enrollment is disabled, the enrollment SAML configuration will be deleted. To enable enrollment, you have to again provide enroll SAML auth server configuration.

8. Confirm that your settings are correct, then select **Create Authentication Server** to create the authentication method.

   The *Authentication Server* page lists the new PingID authentication method.

After you have created your PingID authentication method, create or update your authentication policies with the new authentication method:

1. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > User Policies**.

   The **User Policies** page appears. This page lists all existing user authentication policies.

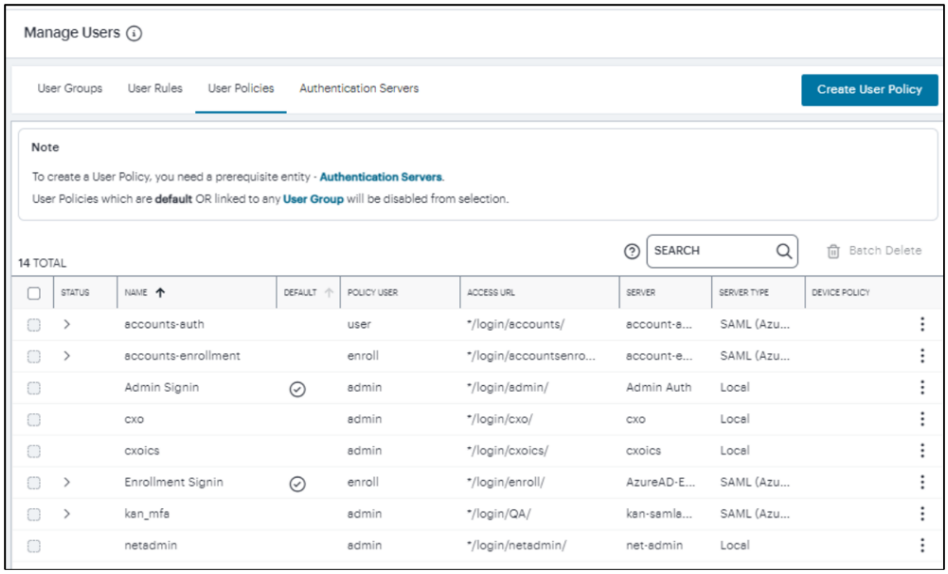

   User Authentication Policies

   To learn more about the policies on this page, see Viewing User Authentication Policies.

   From this page, either create a new custom policy or edit an existing policy.

2. To add a new custom policy, select **Create User Policy**.

   The **Create Authentication Policy** form appears.

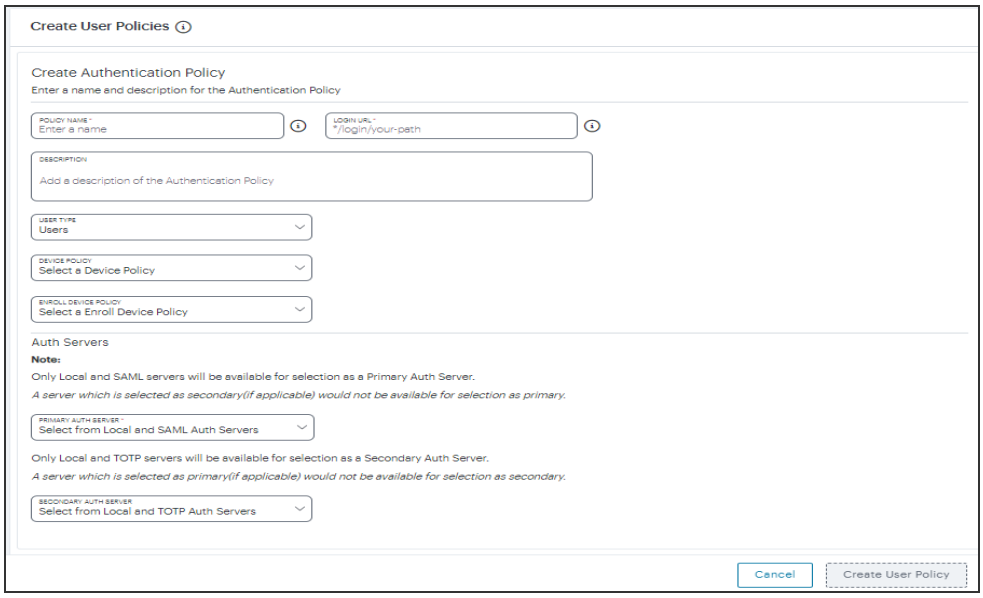

   Create User Authentication

   > At any point during this process, you can reset the form data by selecting **Reset Fields**.

   > To learn more about how custom policies are used for user login and enrollment, see Adding Custom Authentication Policies.

3. Enter a **Policy Name**.

4. Enter a **Login URL** using the format **/login/<path>/*.

   The URL must start with "**/login/*" and cannot contain any special characters. <path> should be set to a unique value reflecting the endpoint URL you want to define for this authentication policy (appended with a backslash):

   - In the case of user sign-in policies, this is the URL endpoint (appended to the tenant FQDN) to which new users are invited to connect to enroll or sign-in a device with the *Controller*. Example value: "**/login/saleslogin/*".
   - In the case of user enrollment policies, this endpoint identifies the enrollment URL to which users are redirected if they attempt to connect to the equivalent sign-in policy with an un-enrolled device. In most cases, you do not advertise this endpoint to your users. Example value: "**/login/salesenroll/*".

     

     In some enrollment circumstances, such as when using a device pre-installed with an older version of *Ivanti Secure Access Client*, you connect directly to the enrollment policy endpoint to enroll the device. For more details, see Using User Authentication Policies.

5. (Optional) Enter a description for the authentication policy.
6. Select a **User Type** based on the intended authentication activity for this policy. Choose from:

   - **Users**: Select this option to define the user sign-in endpoint for enrolled devices. This is the endpoint that you provide to your users to access the service (regardless of enrollment status).
   - **Administrators**: Select this option to define the authentication endpoint for administrator-level sign-in. This endpoint is used for administrator login to the *Controller* only.

7. (for policies with a **User Type** of "Users" only): Select an **Enroll Device Policy** from the drop-down list to be linked to this sign-in policy (as indicated):

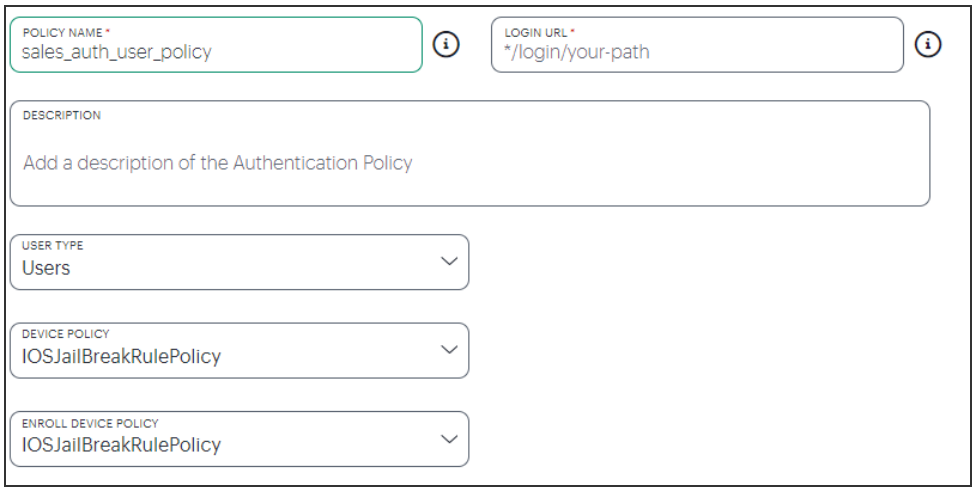

Linking an enrollment policy to a user sign-in policy

This is the enrollment policy to which a user is redirected if it is determined that the device is not yet enrolled. To learn more, see Using User Authentication Policies.

8. Under **Policy Server Details**, select **Primary Auth Server** and choose the required authentication method from the drop-down list:

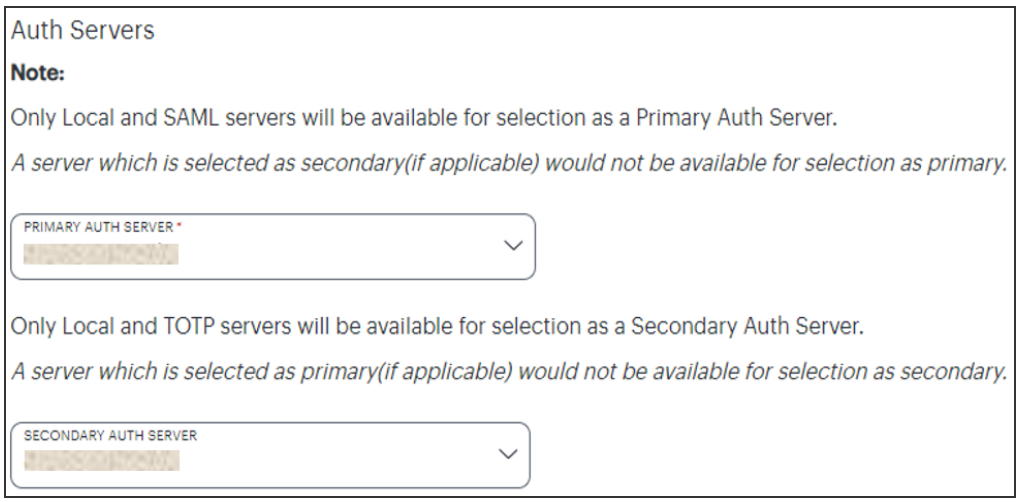

Selecting a primary authentication method for this policy

Alternatively, select *Create Authentication Server* and create a new authentication method as per the steps described earlier in this section.

9. (Optional) Where a secondary method is required for Multi-Factor Authentication, repeat the previous step for **Secondary Auth Server**.

   > Secondary authentication methods do not apply to *Enrollment User* type policies. The relevant field is hidden in this case.

10. Select **Create User Policy** to create the new policy.

    The new policy is added to the list of authentication policies.

If you instead elect to update an existing custom or built-in policy:

1. Click the three dots adjacent to the relevant policy, then select **Edit**.

   The **Update Authentication Policy** form appears.

   > For built-in authentication policies, all properties except **Primary Auth Server** are read-only.

2. Set the **Primary Auth Server** to be the new PingID SAML user authentication method:

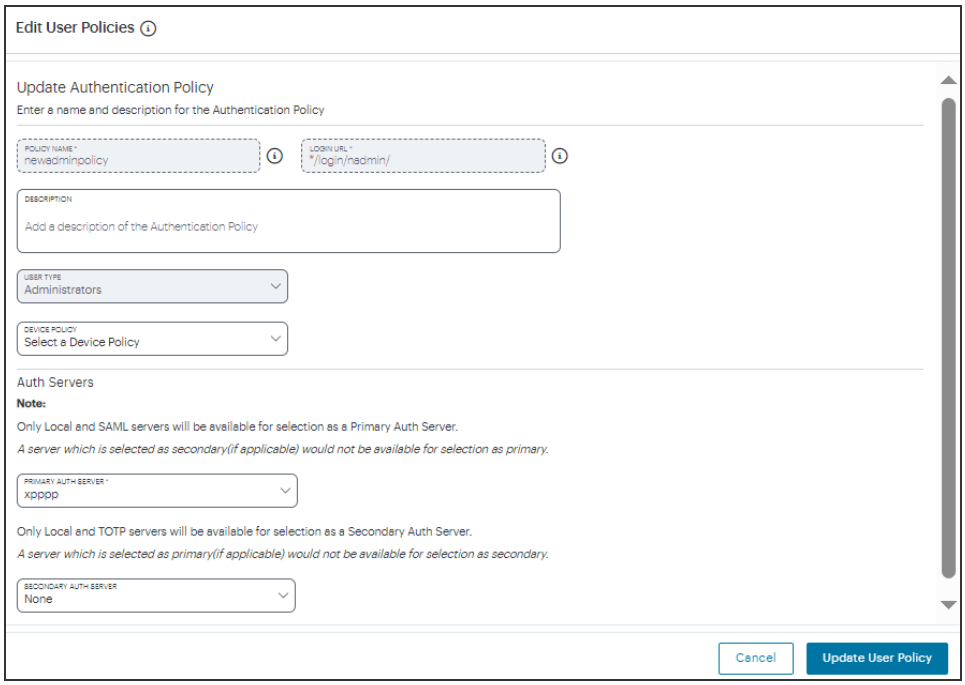

Editing an authentication policy

> SAML authentication can be used only as a Primary Auth Server. If you are using MFA, specify either a *local authentication* or *TOTP* method as the **Secondary Auth Server**.

> If you configure a secondary authentication method in a policy that is currently in use, any active user sessions must be disconnected and reconnected for the changes to take effect.

3. Select **Update User Policy**.

   The list of authentication policies updates.

On the **Authentication Policies** page, select the authentication policy you just created (or updated), then select the corresponding **Download** link. Save the resulting SP metadata file to your local workstation.

After you have completed this process for both enrollment and sign-in activities, you can proceed to finish configuring your PingID application using the SP metadata file downloaded during this workflow. See Updating Your PingID SAML Application with Your SP Metadata.

Updating Your PingID SAML Application with Your SP Metadata

> To fully configure PingID as a SAML authenticator in *nZTA*, complete these steps for each of your user enrollment and user sign-in policies in turn.

After you have obtained the SP Metadata file from your *nZTA* Authentication Policy, update the PingID SAML application created in Creating a PingID SAML Application. Make sure you use the application that corresponds to the policy. In other words, if you have obtained metadata for the enrollment policy, update the PingID enrollment application.

To update the PingID application:

1. Return to the PingID application previously started through the PingOne portal (https://admin.pingone.com):

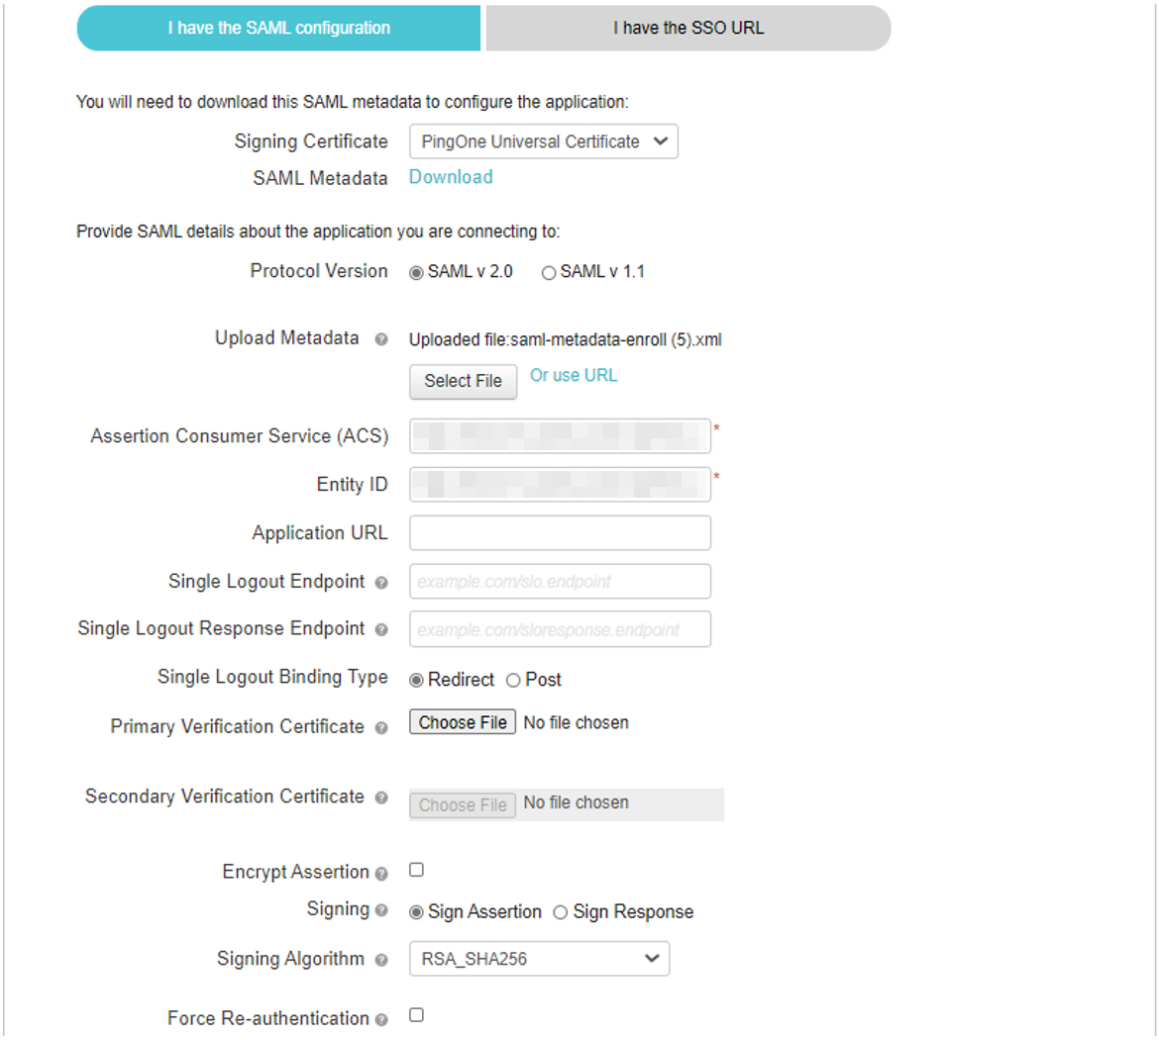

Continuing the PingID application creation process

2. Locate the **Upload Metadata** field and select **Select File**. Upload your SP metadata file.

3. Make any further changes you require to the application configuration, then select **Next**.

   The **Group Access** section appears.

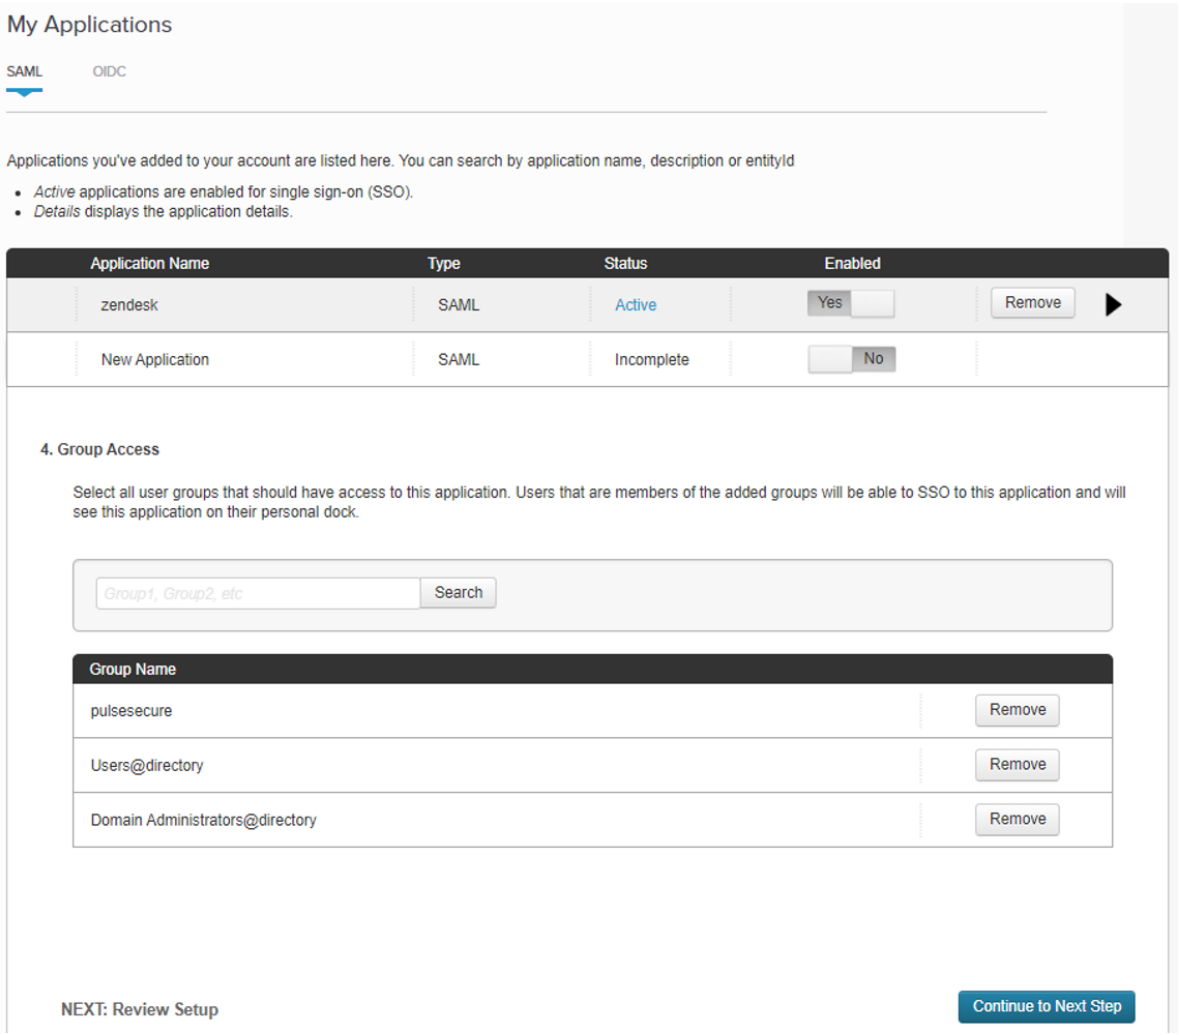

   Adding groups to your PingID application

4. Add to your application any required user groups.

> The user groups you require depend on your organization and might vary from those shown in this workflow. At least one user group is required.

5. To review your application configuration, select **Continue to next step**.

   The **Review Setup** stage appears.

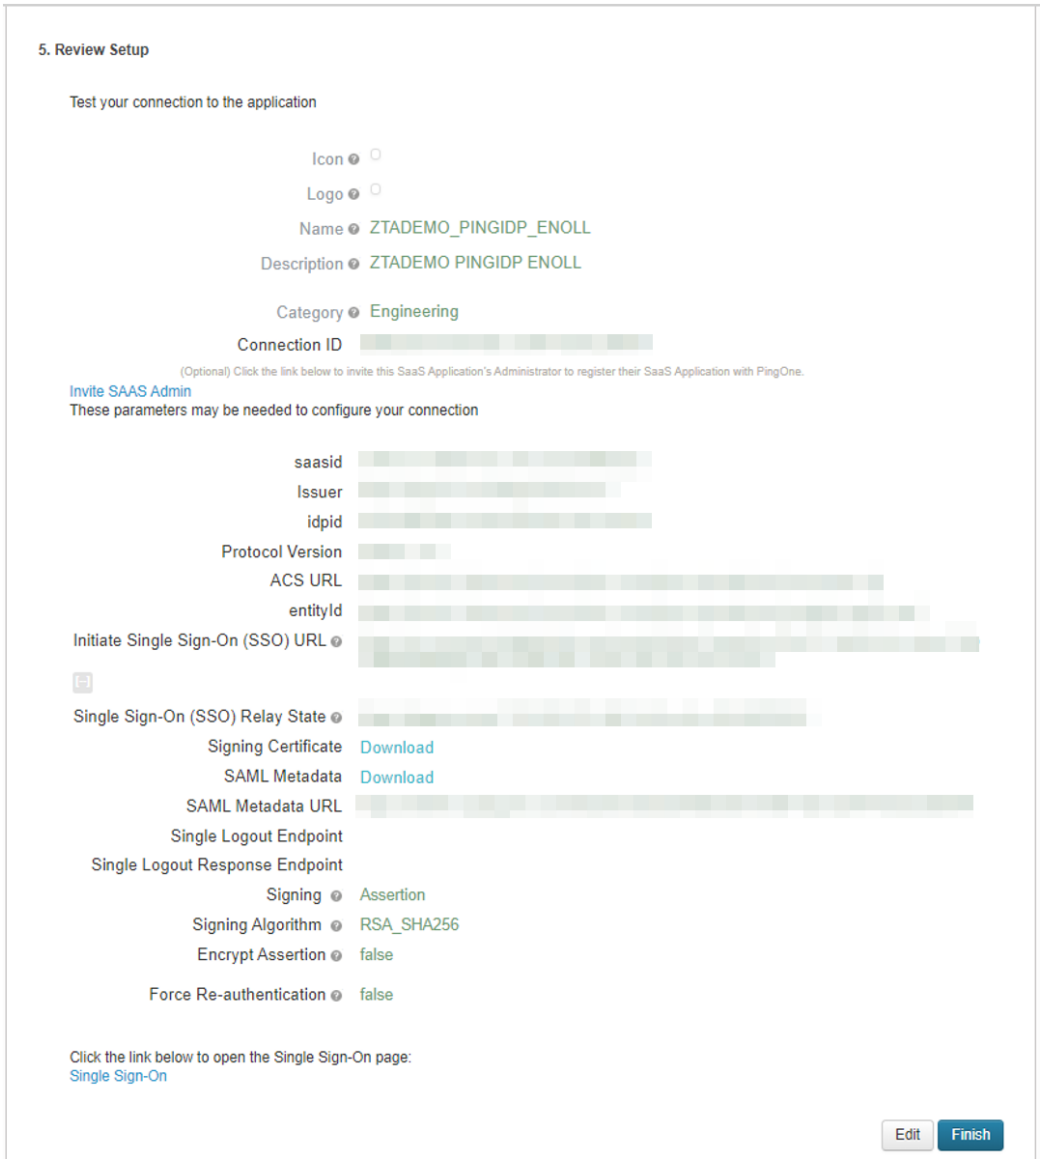

   Reviewing your application configuration

6. To complete configuration of your PingID application, select **Finish**.

Workflow: Adding TOTP to an Authentication Policy

> This feature is supported for client and gateway versions applicable to release 22.2R1 and later only.

nZTA supports the use of Time-based One Time Password (TOTP) as a secondary authentication method in *Multi-Factor Authentication* deployments.

To use TOTP, first create a TOTP authentication method in *nZTA* and then associate it with your user sign-in authentication policies.

To ensure that your users can access the authentication mechanism defined in the policies you configure through this process, make sure your *Secure Access Policies* are configured with a **User Group** in which these authentication policies are defined. To learn more, see Creating/Editing Secure Access Policies.

To configure a new TOTP authentication method:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > Authentication Servers**.

   The **Authentication Servers** page appears. This page lists all existing user authentication methods. For example:

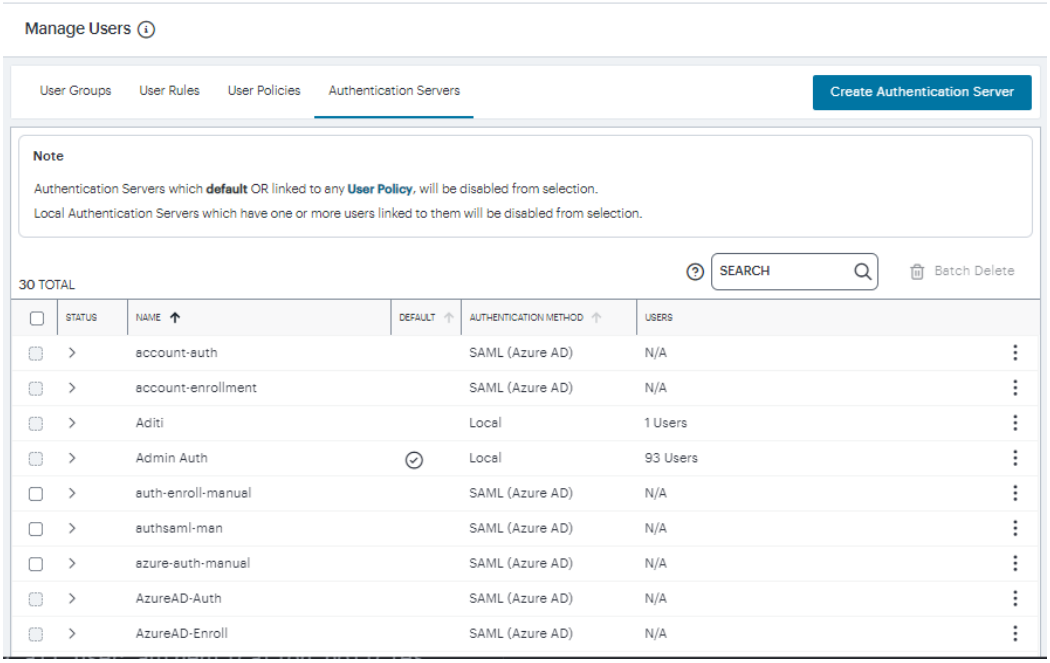

User Authentication Methods

3. Select **Create Authentication Server**.

   A form appears that enables you to define the authentication method.

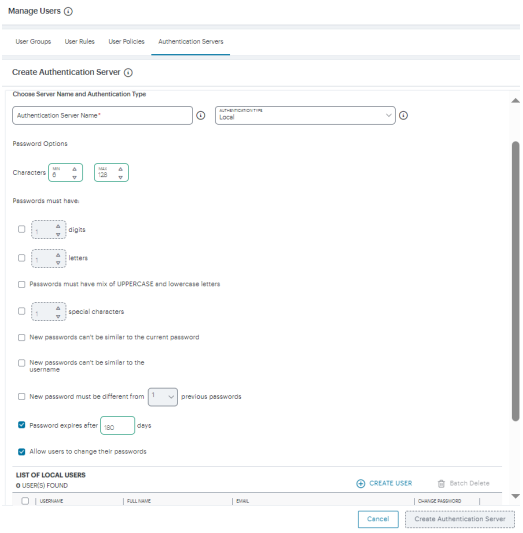

Creating a new TOTP user authentication method

> At any point during this process, you can reset the form data by selecting **Reset Fields**.

4. Under **Choose name and type**:

    - Specify an **Authentication Server Name**.
    - Select the **Authorization Type** of *TOTP*.

    The form expands to show additional TOTP authentication settings:

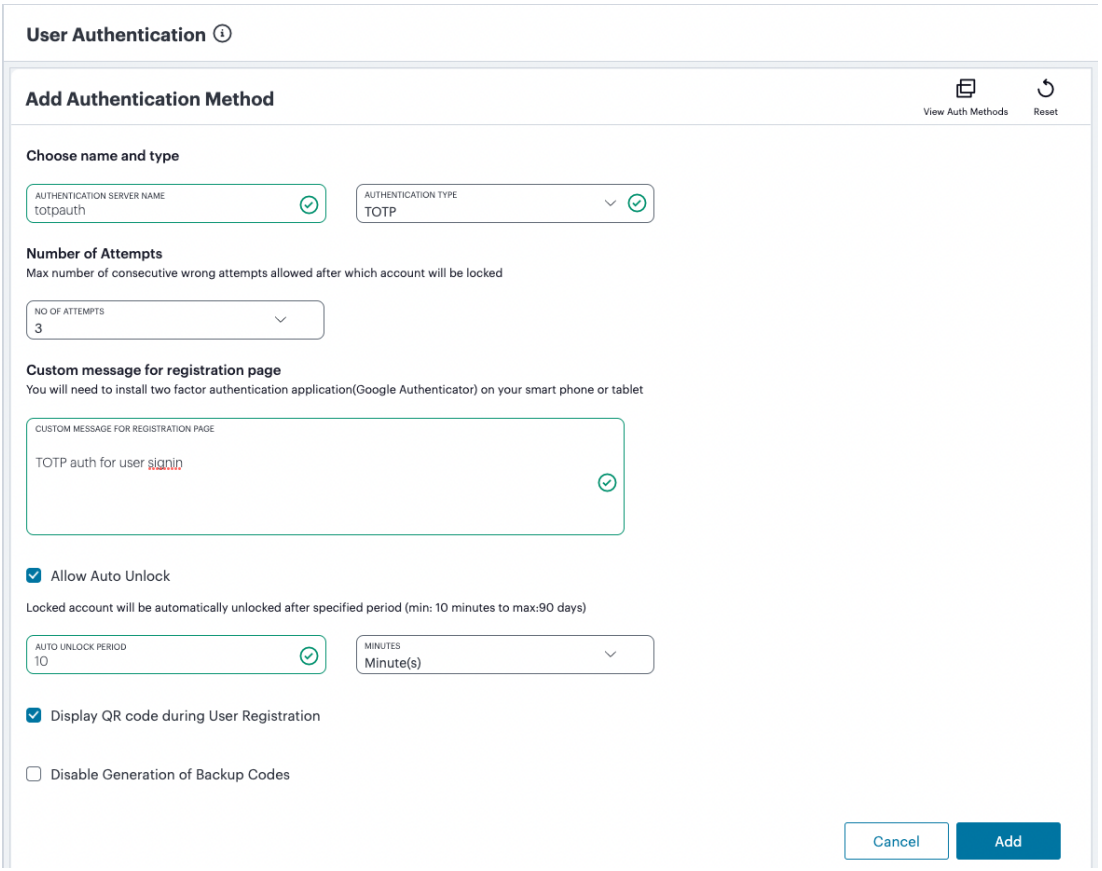

Adding TOTP authentication settings

5. Enter the following settings:

    - **No of Attempts**: The maximum number of consecutive wrong attempts allowed before the account is locked (minimum: 1 attempt, maximum: 5 attempts). To view user attempts and to unlock locked accounts, see Unlocking Locked User Accounts.
    - **Custom message for registration page**: A custom message to be shown on the new TOTP-user registration web page.
    - **Allow Auto Unlock**: When selected, a locked account is automatically unlocked after the specified period. (minimum: 10 minutes, maximum: 90 days).
    - **Display QR code during User Registration**: When selected, a QR code is displayed during user registration.
    - **Disable Generation of Backup Codes**: When selected, the *Controller* does not generate TOTP backup codes.

6. To create an authentication method based on these settings, select **Add**.

    The new TOTP user authentication method is added to the list of methods and the process is complete.

After you have created your TOTP authentication method, create or update your user sign-in authentication policies with the new method. *nZTA* supports using TOTP *only as secondary authentication*, so make sure you have previously configured a primary authentication method before continuing this process. To view workflows for all available authentication types, see Introduction.

Secondary authentication methods do not apply to *Enrollment User* type policies. The relevant field is hidden in this case.

Complete the following steps for your user sign-in authentication policy:

1. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Users > User Policies**.

   The **User Policies** page appears. This page lists all existing user authentication policies.

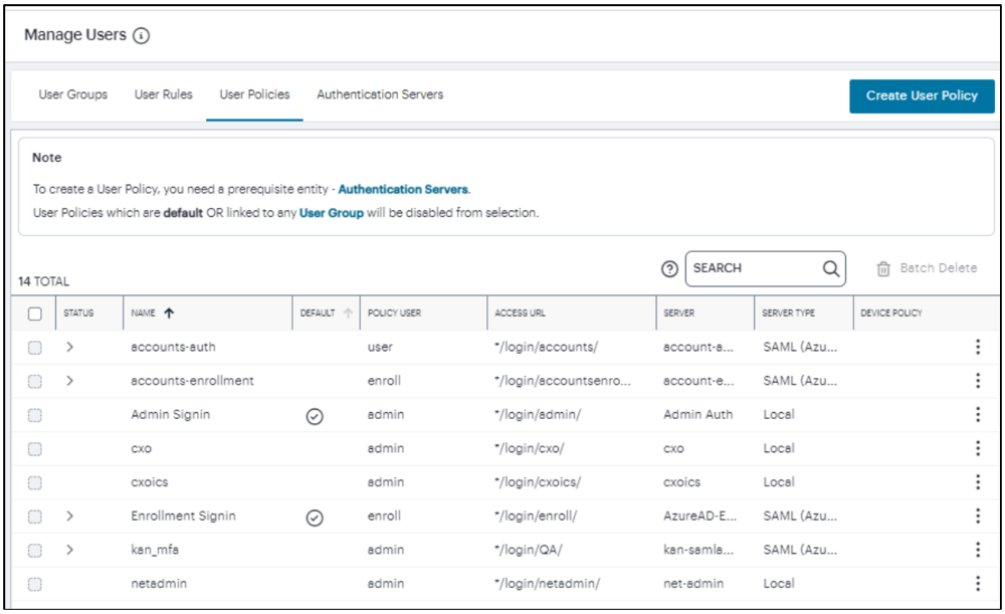

   User Authentication Policies

   To learn more about the policies on this page, see Viewing User Authentication Policies.

2. Click the three dots adjacent to your desired user sign-in policy, then select **Edit**.

   The **Edit Authentication Policy** form appears.

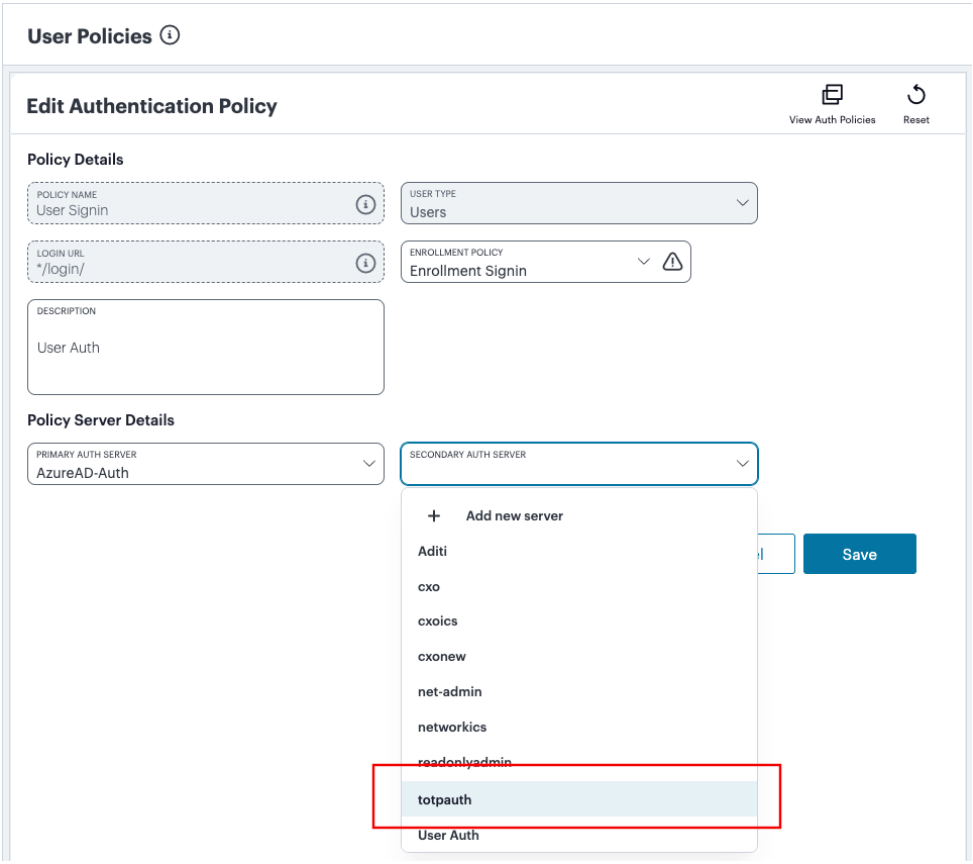

Selecting a secondary TOTP authentication method for this policy

3. For **Secondary Auth Server**, select your new TOTP authentication method from the drop-down list (as indicated).

   Alternatively, select *Add New Server* and create a new authentication method as per the steps described earlier in this section.

4. Select **Save** to update the policy.

> If you configure a secondary authentication method in a policy that is currently in use, any active user sessions must be disconnected and reconnected for the changes to take effect.

Unlocking Locked User Accounts

After you have created a TOTP authentication method and assigned it to an active user authentication policy, you can use the authentication method configuration page to view users that have attempted authentication through TOTP. This information enables you to unlock locked user accounts, if required.

To access user attempt information, perform the following steps:

1. Select **Secure Access > Manage Users > Authentication Servers**.
2. Click the three dots adjacent to your TOTP authentication method, then select **Edit**.

   At the bottom of the page, a Users table is presented:

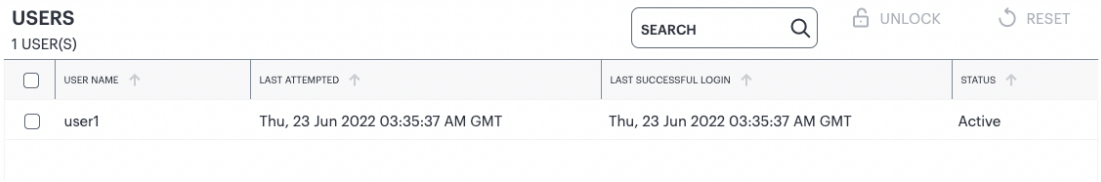

Viewing the list of users who attempted TOTP authentication through this method

   This table lists each user who has attempted to authenticate a device through TOTP, including the last attempt and last successful login times.

3. (Optional) If a user account is locked through too many consecutive failed authentication attempts (that exceed the value configured in **No of Attempts**), unlock the account by selecting the checkbox adjacent to the user entry and selecting **UNLOCK**. The user is then free to re-attempt authentication using valid authentication codes.
4. (Optional) To remove a user from the list, select the checkbox adjacent to the user entry and select **RESET**. This means a user must then re-register their device with the TOTP policy.

Reset and unlock operations of individual users are supported only when the TOTP authentication method is associated with a user authentication policy. To reset or unlock all users in a disassociated TOTP authentication method, delete the TOTP authentication method itself.

Working with Gateways

Introduction

After you have successfully logged into the *Ivanti Neurons for Zero Trust Access* (*nZTA*) *Controller* as a Tenant Admin user (see Logging in as a Tenant Administrator), you can start the configuration of your *nZTA* platform by adding Gateways.

nZTA supports two main Gateway types, depending on your subscription:

- *nZTA Gateway*
- *Ivanti Connect Secure* (*ICS*) Gateway

This chapter covers functionality relating to a *nZTA Gateway* only. For details pertaining to a *ICS* Gateway, see instead the "*ICS* Tenant Admin Guide" in the *nZTA* documentation portal.

To learn more about *nZTA Gateway*s and their relationship with the other dimensions of a *Secure Access Policy*, see Deploying Gateways.

A Gateway controls access to the applications at the location to which it is deployed. This location could be a physical datacenter, a private or public cloud-based service, or some hybrid combination. Each Gateway communicates with the *Controller* to ensure that access requests are authenticated. A Gateway must be contactable by both the *Controller* and the applications that reside there. See Configuring Networks in your Gateway Datacenter.

The process of defining a Gateway on the *Controller* produces a package of settings, known as a Gateway definition, that you publish to the Gateway virtual machine instance during deployment. These settings enable the Gateway to establish communication with the *Controller*.

Ensure the Gateway virtual machine instance does not exist prior to creating your Gateway definitions on the *Controller*. All *nZTA Gateway*s must be deployed from the *Controller* directly. The Gateway definition file is designed to be published to a new virtual machine Gateway instance during its initial deployment.

To register a new Gateway with the *Controller*, use the Tenant Admin portal. The *Controller* requires basic identification and networking details for the Gateway instance, and in return provides a downloadable Gateway definition file.

A Gateway definition file is valid for 24 hours. If this period expires, you must replace the Gateway to generate a new Gateway definition file.

When you deploy the Gateway virtual machine instance in your on-premise or cloud infrastructure, you apply the definition file as configuration data. At launch time, the Gateway attempts to contact and register itself with the *Controller*, establishing a link between the Gateway record in the *Controller* and the actual virtual machine instance. Any subsequent policy changes made on the *Controller* are automatically synchronized out to all Gateways.

You can choose a single Gateway (or Gateway Group) to act as a *default Gateway*. A default Gateway handles all requests from applications that are not referenced by any secure access policy. See Viewing and Monitoring Gateways in the Controller.

> *Ivanti Secure Access Client* Linux variants do no currently support the use of a default gateway.

After you have registered your Gateways with the *Controller*, you can:

- View performance and usage graphs
- See activity logs
- View task lists
- Establish Gateway Groups for *High Availability*
- Manage version upgrades
- Replace Gateway instances registered with the *Controller*

To learn more, see Viewing and Monitoring Gateways in the Controller.

High Availability

nZTA allows you to deploy multiple Gateways in front of the same set of applications or resources to support high availability. This arrangement can be used to provide scaling, redundancy, and load distribution for your application delivery. To learn more, see Using Gateway Groups for High Availability.

High availability is implemented in the *Controller* through **Gateway Groups**. You add individual Gateways to a group, and then associate the group with your Secure Access Policy. To learn more about adding Gateway Groups, see Adding Gateway Groups for High Availability.

Gateway Deployment Workflows

nZTA supports Gateway virtual machine instances deployed in the following environments:

- **VMware vSphere**: see Workflow: Creating a Gateway in VMware vSphere.
- **Amazon Web Services (AWS)**: see Workflow: Creating a Gateway in Amazon Web Services.
- **Microsoft Azure**: see Workflow: Creating a Gateway in Microsoft Azure.
- **KVM on OpenStack**: see Workflow: Creating a Gateway in KVM/OpenStack.
- **Google Cloud Platform**: seeWorkflow: Creating a Gateway in Google Cloud Platform.

You can choose any single Gateway (at v21.1 or later) or Gateway Group to act as a *default Gateway*. A default Gateway handles all requests from applications that are not referenced by any secure access policy. See Configuring a Default Gateway for Application Discovery.

> Each process described in this chapter contains prerequisites that correspond to the latest supported Gateway versions for this release. To deploy older supported Gateway versions, substitute in the file paths and names with the version you want to use. To learn more about the supported Gateway images for this release, see the Release Notes or contact your support representative.

Configuring Networks in your Gateway Datacenter

*nZTA Gateway*s deployed in your cloud or on-premise datacenter require the availability of a number of network interfaces and ports to operate correctly. You require the following primary interfaces defined in your Gateway virtual machine instance:

- **External network interface**: Configured with a public subnet IP address and used for external client access to the applications deployed in that datacenter. Use this IP address during the process of creating your Gateway record on the *Controller*.
- **Internal network interface**: Configured with a private subnet IP address and used for internal connections to the deployed applications, and for external communication with the *Controller*.
- **(Optional) Management network interface**: Configured with an IP address and port on a further, separate, network subnet for deployments where a specific management interface is required.

> When the management port is enabled, Gateway will use management interface to communicate with Controller and NTP Server. The Gateway will still use the internal port for DNS resolution and NTP server name resolution. If the internal DNS cannot resolve the Controller domain, the internal interface will require internet access.

nZTA Gateway template images contain parameters and settings for all required interfaces. You provide suitable configuration during the Gateway deployment workflows described in this chapter.

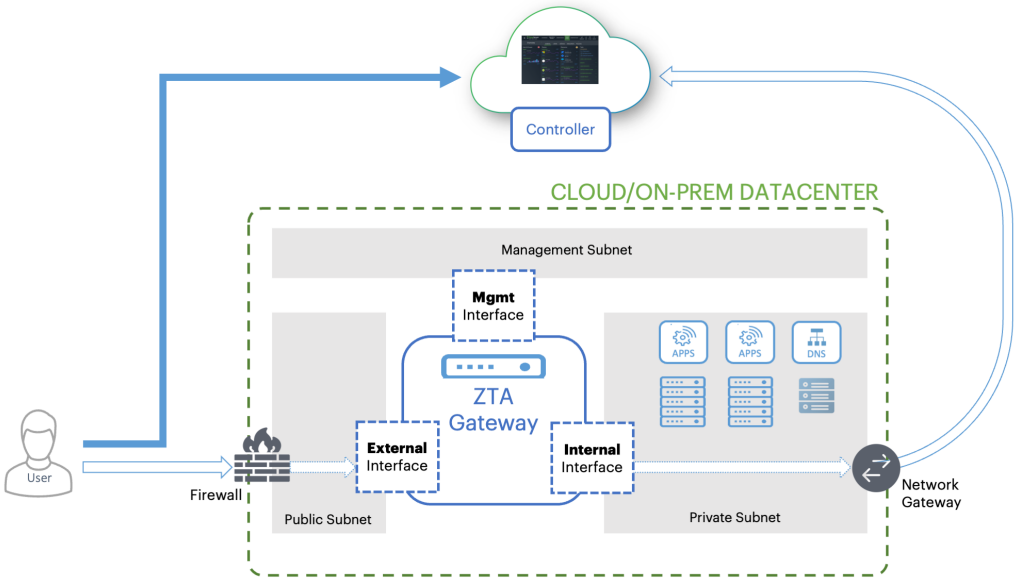

Gateway network connections in your cloud and on-premise datacenter

nZTA Gateway deployment using this topology offers the following attack protection benefits:

- **mTLS Communication**: By using mTLS for communication with the *Controller*, your Gateways are protected from traffic originating from unauthorized and un-enrolled devices.
- **DMZ/DDoS Deployment**: *nZTA* allows customers to deploy their own DDoS (Distributed Denial of Service) mitigations, either through network firewalls or DDoS services, before traffic reaches the *nZTA Gateway*s deployed in your datacenter.
- **SSL inspection**: If an administrator chooses to use SSL interception or deep inspection for traffic coming into the network or to back-end applications, this can be achieved on the internal *Private Subnet* side before traffic reaches your back-end application servers.

The *Controller* communicates with all registered Gateways to validate user sessions, and for reporting/analytics. Each Gateway maintains a timeout period of 5 minutes for validation of user sessions. If a Gateway fails to obtain a response from the *Controller* within this time limit, any new session authorization requests are denied. However, existing sessions remain connected to the applications and resources at that location until the user session expires. To learn more about user session expiry, see On-Demand and Simultaneous Connection Handling.

For all platforms, make sure the firewall rules for the **Public Subnet** in which your *nZTA Gateway***External Interface** resides is configured to accept inbound client connections on TCP port 443.

Furthermore, make sure you configure the **Network Gateway** serving your **Private Subnet** to allow outbound traffic to the *Controller* in the following ways:

- Allow outbound TCP traffic on port 443 to the *Controller* service
- Allow outbound UDP traffic to the following Network Time Protocol (NTP) services:
    - time.windows.com (port 123)
    - time.nist.gov (port 123)

If you maintain your own DNS service at the datacenter, you can specify these details during Gateway deployment.

If you are planning to use your *nZTA Gateway* to serve SaaS (Software-as-a-Service) applications, configure your application to restrict inbound connections to your **network gateway** IP address. This ensures that your SaaS application can be reached only by clients connecting through the *nZTA Gateway*.

Using Dynamic IP Addressing to Profile Client Traffic

This feature is supported for VMware (ESXi) and KVM Gateway types only.

This feature is not supported for Gateway Groups.

A client establishes a secure tunnel to a Gateway in order to reach the applications and resources controlled by a Secure Access Policy. Traffic from the client passes through the tunnel to the Gateway, and from the *Internal* network interface on the Gateway to the application. The application or resource sees the client's traffic as originating from the Gateway's internal interface. This scenario is transparent to the client and application - the Gateway manages traffic back to the client using Network Address Translation (NAT) to map traffic on the internal interface to the client's source IP address.

However, in some circumstances you might want to profile or monitor end-to-end traffic for your clients. This is difficult beyond the Gateway as traffic at the application appears to originate from the same source IP address (that of the Gateway internal interface). To facilitate this, *nZTA* includes the option to specify a pool of IP addresses in a dedicated subnet that the *Controller* can dynamically assign to client sessions. As a client sends traffic to an application or resource, the Gateway establishes a mapping from the tunnel IP address to one of the free IP addresses in the pool. This dynamically-assigned IP address is then used as the source IP address when sending client traffic to the application. The Gateway again uses NAT to manage the connections to each client.

You can configure dynamic IP addressing when adding a new Gateway or by editing an existing Gateway configuration. To learn more, refer to the Gateway configuration workflows described in this chapter. To read more about how to enable Dynamic IP Addressing in an existing Gateway, see On-Demand and Simultaneous Connection Handling.

To use dynamic IP addressing, the *Controller* requires you to define a unique address range for each applicable Gateway, using CIDR notation. For example, `192.0.2.0/24`. You can add only one IP address range per Gateway.

> The allowed subnet range is 8-28. Make sure you select a subnet value that provides the amount of IP addresses necessary to map the expected number of clients connecting to the Gateway. If you exhaust the IP address range, your clients can still connect, although traffic profiling is not possible.

When configuring an address range, make sure this does not overlap with dynamic IP address ranges assigned to other Gateways.

White-listing Required IP Addresses for your Services

The *Controller* service uses a series of IP addresses and ports to facilitate access to the admin and user web consoles, for user enrollment, and for connections to *nZTA Gateway*s. To ensure network access, make sure the following IP addresses and ports are white-listed (or added to the *allowed list*) in your network firewalls and routing infrastructure.

Select the IP addresses and ports for your corresponding region only:

- **North America**:

  52.186.44.249 (port 443)

  52.188.33.186 (port 443)

- **Europe**:

  51.138.111.17 (port 443)

  20.50.150.82 (port 443)

- **APJ**:

  20.44.238.229 (port 443)

  20.44.237.67 (port 443)

- **UAE**:

  20.233.40.108 (port 443)

  20.233.41.69 (port 443)

- **Canada**:

  20.220.157.85 (port 443)

  20.220.157.158 (port 443)

Viewing and Monitoring Gateways in the *Controller*

To view, configure, and monitor the health of your deployed Gateways and Gateway Groups, use the **Secure Access > Manage Gateways** section of the *Controller* Tenant Admin portal. The pages in this section remain inactive until you select a Gateway or Gateway Group.

> To view detailed analytics for your deployed Gateways, see also Monitoring nZTA Gateway Activity.

To view information for a Gateways or Gateway Group:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways > Gateways List**.

   The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*:

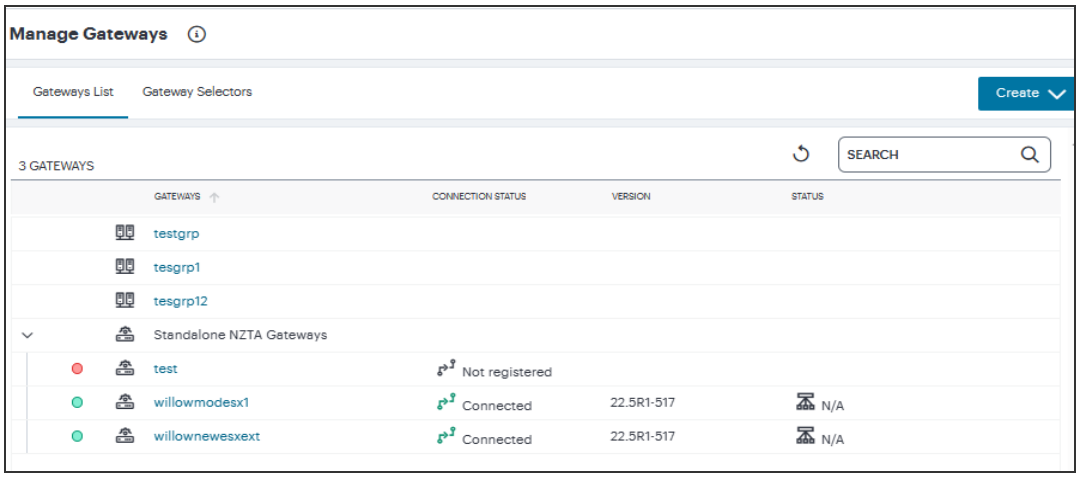

   Viewing All Gateways

The health and status of each Gateway instance is denoted by the indicator colors used in the left-most column; *green* for connected/ready or *red* for disconnected/not ready. Unregistered Gateways are denoted by a missing indicator. Additional status information is given under *Connection Status*, with the current software *Version*, and the *Status* of any scheduled tasks.

Use the *Search* box at the top enter a text search term to narrow the list to only matching Gateways and Gateway Groups.

3. Select a Gateway or Gateway Group from the list to view the *Gateways Overview* page.

   This page provides further status and configuration details for the selected instance, as well as *Performance*, *Concurrent Sessions*, and *Throughput* usage graphs.

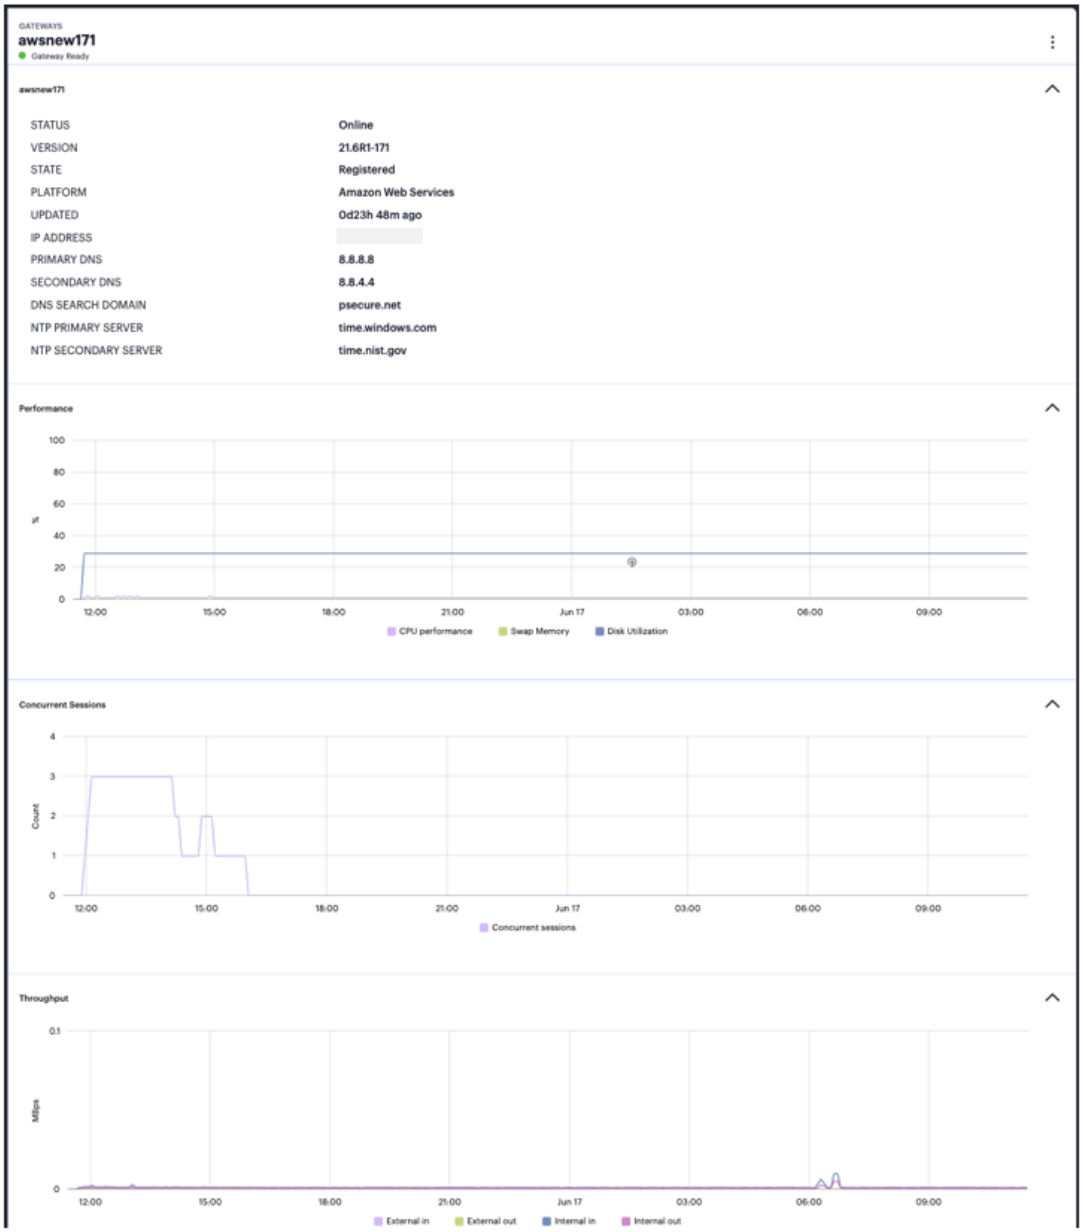

   Viewing status, configuration, performance and Usage Graphs in the Gateway Overview page

   Hover your pointer over a coordinate in each graph to view a tooltip showing detailed metrics for that moment.

For Gateway Groups:

- **Edit Group**: Edit details for the group, such as name, description, and load balancer public IP address.
- **Delete Group Only**: Deletes the Gateway Group record, leaving any contained Gateways intact as standalone (ungrouped) instances.
- **Delete Group and Gateways**: Deletes the Gateway Group and all Gateways contained in the group.

This option appears only if the selected Gateway is in a Gateway Group.

- **Add Gateway to this Group**: Add a registered standalone gateway to this group.

For Gateways:

- **Edit Gateway**: Edit details for the Gateway, such as name, description, and public IP address. You can also select or reset the **Use Management Port** property.
- **Replace Gateway**: Allows replacement of the Gateway virtual machine instance associated with this *nZTA Gateway* record. This option regenerates a Gateway definition file based on the existing networking details stored in the Gateway record, but with a new one-time token, to allow deployment and registration of a new virtual machine instance.

A Gateway virtual machine deployed to replace an existing Gateway might be launched with a different public IP address on the external network interface. To update the public IP address setting stored for the Gateway in *nZTA*, use the *Gateway Network Settings* section of the **Secure Access > Manage Gateways > *Gateway* > Configuration** page. For more details, see Editing Gateway Configuration.

- **Remove from Group**: Removes this Gateway from a Gateway Group.

This option appears only if there are existing Gateways in the Group.

- **Delete Gateway**: Deletes the Gateway record.
- **Upgrade to <version>**: Upgrades the registered Gateway virtual machine instance to the specified version, see Upgrading Gateways.
- **Rollback to <version>**: Reverts the registered Gateway virtual machine instance to the specified version, see Upgrading Gateways.

The available menu options vary depending on whether the selected Gateway is registered and connected to the *Controller*.

Viewing Gateway Logs

The *Logs* page enables you to view the Access, Admin, and Event logs for the selected Gateway or Gateway Group.

To view Gateway logs:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways > Gateway List**.

   The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*.
3. Select the required Gateway or Gateway Group.

   The *Gateways Overview* page appears for the selected Gateway/Gateway Group.
4. Select the **Logs** tab.

   The *Gateways Logs* page appears.

This page supports the following features:

- Select the log type you want to view in the **Log Type** drop-down list. Choose from:
  - Access Logs
  - Admin Logs
  - Event Logs

- Set a **Time Period** over which the logs are shown. Use the time period selector to set a time period or range for your log results. Click the calendar (highlighted) to show the selector:

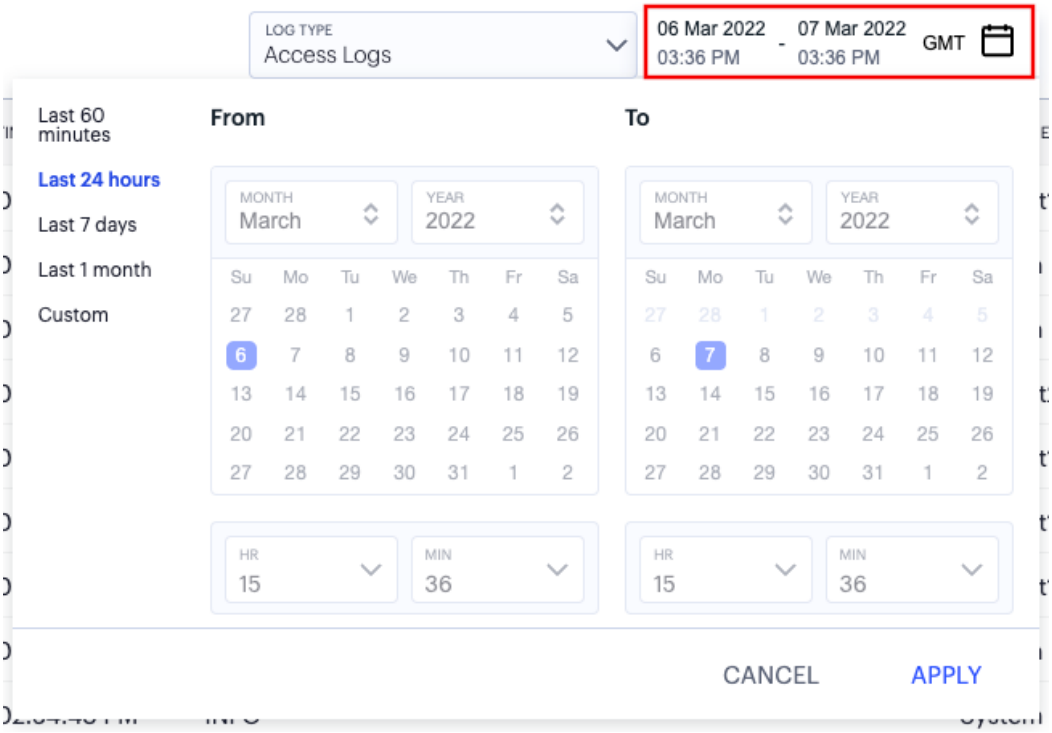

Setting a log time period

Set the time period you want to view using the available ranges at the top-left. Choose from:

- Last 60 minutes
- Last 24 hours (default)
- Last 7 days
- Last 1 month
- Custom

For **Custom**, set a specific *From* and *To* to denote the start and end of your custom date/time range.

> The custom date/time calendar controls are enabled for only the **Custom** option. However, the calendar continues to identify the applicable start and end date-time for all predefined time periods.

To apply your changes, click **Apply**. The selected time period is displayed in the filter bar and data on the page updates accordingly.

> To configure the timezone, see Setting the Timezone.

- Logs are refreshed automatically by changing the criteria. To manually refresh the log display, click the following icon:

  

  Refreshing the page data

- To change the fields displayed for each log line, use the following icon:

  

  Show or hide log fields

  In the field selector, click a field name to toggle between show or hide. A *tick* icon indicates a displayed field. After you are finished, click the context menu icon to close the selector.

  Choose from the following fields:

  - Session identifier
  - Gateway identifier
  - Gateway name
  - Source IP address
  - User name associated with the event, where applicable
  - User device identifier, if available
  - A description of the event

- Use the following icon to trigger the advanced filter selection:

  

  Applying a filter to the log display

  Use the filter to narrow down the displayed log records to your selected criteria. For example, to show only those log messages with a critical severity level, or pertaining to a specific user, or both. To learn about log filtering, see Filtering the Logs.

- Use the following field to use search term highlighting. Enter a value into the search box, *nZTA* highlights all matches in the log display.

  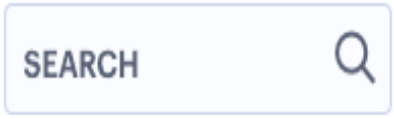

  Highlighting a search term in the logs

- To switch between the default and denser data views, use the following icon:

  

  Setting the view density

- To apply grouping to the displayed log records, click the following icon:

  

  Group log records by selected criteria

  This feature applies grouping to a selected field in the log record display, such that records are accumulated and grouped together under each unique data item identified in that field. Through grouping, an admin can quickly view the number of records of a particular type.

  To learn more about record grouping, see Viewing Detailed Logs for a Chart.

- To view a single log entry in a dedicated panel, click the log message text to activate the info-panel view:

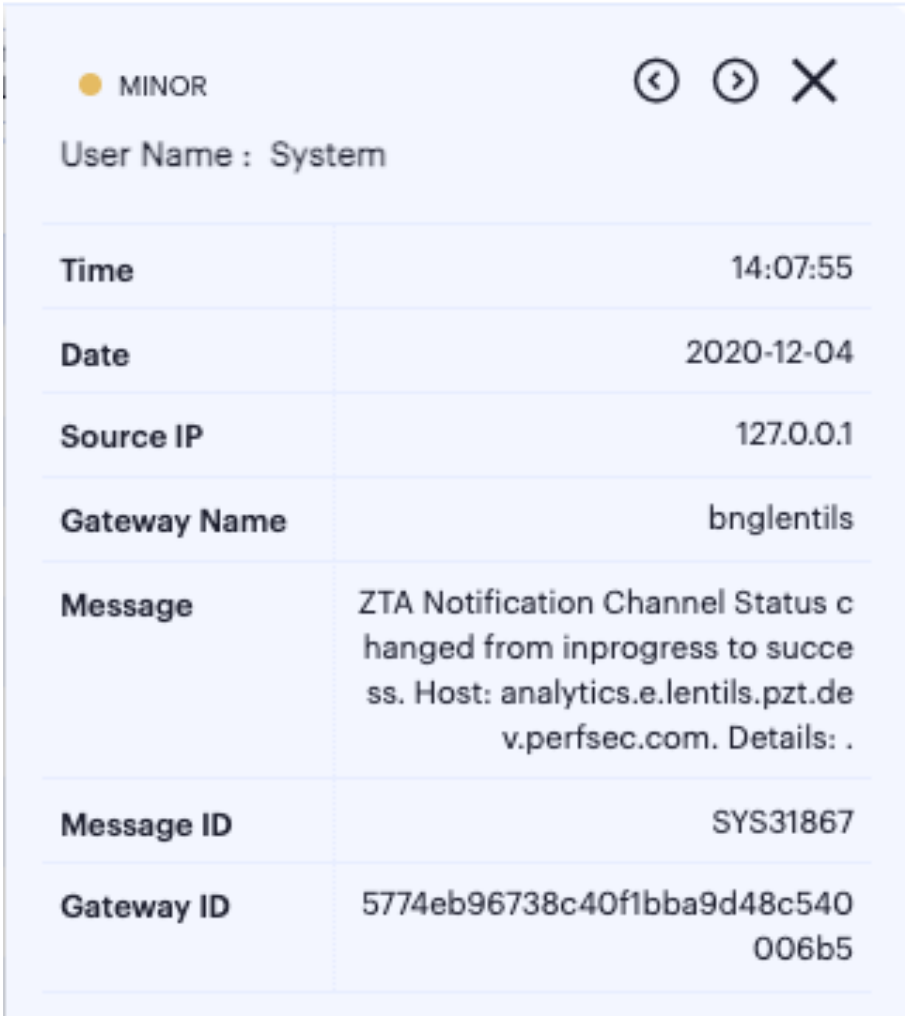

Viewing a single log entry in the info-panel

In the info-panel, use the arrow icons to cycle through the previous and next log entries in turn.

- To view, sort, and filter the log messages in context with all other logs in your deployment, click **VIEW IN LOGS PAGE** (see Checking the Logs). This page supports the following features:

Viewing Gateway Tasks

The *Tasks* page enables you to view the current task list for the selected Gateway or Gateway Group. A task is triggered when an action on a Gateway requires an update to be made to the Gateway instance. For example, if you add a Gateway to a Gateway Group, two tasks are created: A Gateway Group change, and a change of certificate on the joining Gateway.

This list is read-only and cannot be modified. To filter the displayed tasks, use the **Task Type** and **Date** controls at the top of the page. Task type contains the primary categories within which each task falls. To view all tasks for the current day, click **Clear**.

To view Gateway tasks:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways > Gateway List**.

   The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*.
3. Select the required Gateway or Gateway Group.

   The *Gateways Overview* page appears for the selected Gateway/Gateway Group.
4. Select the **Tasks** tab.

   The *Gateway Tasks* page appears, showing the task list for the selected Gateway/Gateway Group.
5. (Optional) Select the **Task Type** and **Date** to filter the results.

Editing Gateway Configuration

The *Configuration* page enables you to edit the selected Gateway or Gateway Group configuration.

To edit a Gateway/Gateway Group:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways > Gateway List**.

   The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*.
3. Select the required Gateway or Gateway Group.

   The *Gateways Overview* page appears for the selected Gateway/Gateway Group.
4. Select the **Configuration** tab, or click *Edit Gateway/Group* from the context menu at the top-right.

The *Gateway Configuration* page appears, showing the current settings for the selected Gateway/Gateway Group:

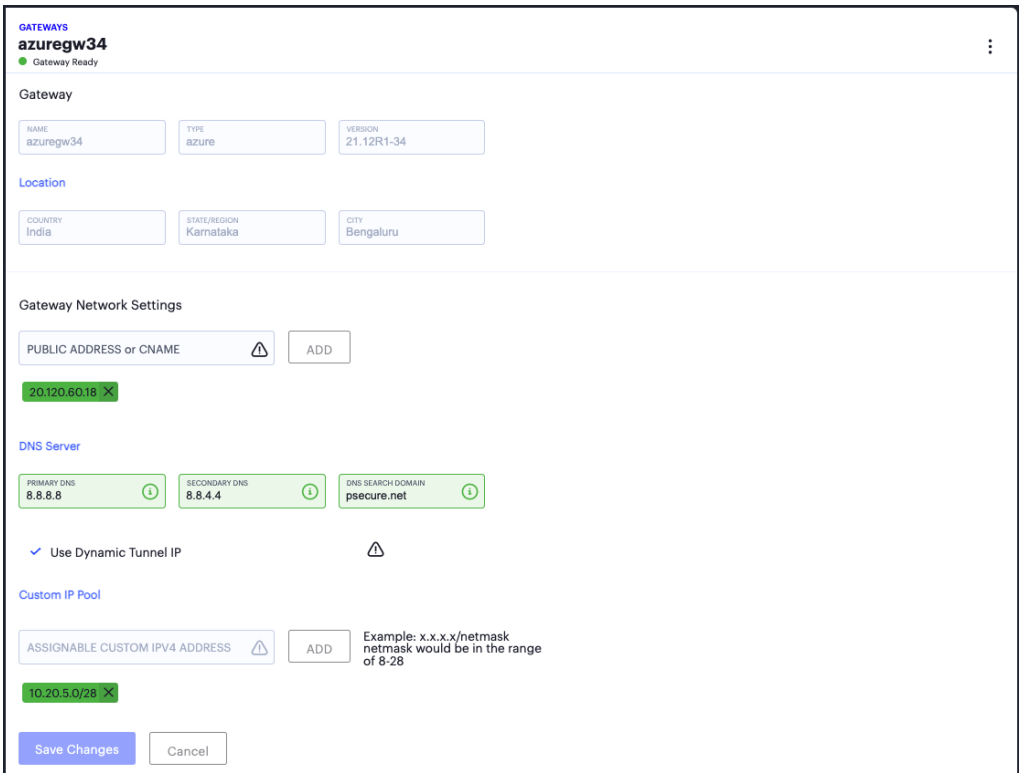

Editing a Gateway configuration

On this page, you can edit the following settings:

Editable Gateway Network Settings

Setting	Description
Public Address or CNAME	The public IP address or CNAME at which clients can externally reach the Gateway instance. To learn more, see Configuring Networks in your Gateway Datacenter.
DNS Server	Contains settings relating to the Domain Name Service (DNS). Enter your **Primary DNS** and **Secondary DNS** IP addresses, and **DNS Search Domain**.
(vSphere Gateways only) Use Manual Settings	Allows you to manually configure IP address settings for your Gateway **Internal**, **External**, and (optional) **Management** interfaces. Deselect this option to instead use DHCP.

Setting	Description
(vSphere and KVM Gateways only) Use Dynamic Tunnel IP	Allows you to assign a pool of IP addresses that are dynamically mapped to client sessions, such that user traffic from the Gateway to the application can be identified as originating from a specific client. Enter an IP address and subnet (in the range 8-28) in CIDR notation, then click **Add**. To learn more about , see Dynamic IP Addressing to Profile Client Traffic.

5. Update any required Gateway/Gateway Group details, then click **Save Changes**.

Troubleshooting Gateway Issues

The *Troublshooting* page provides tools that enable you to investigate issues that might be affecting your Gateway or preventing it from operating normally.

> Troubleshooting is available only for fully registered Gateways, and is not supported for Gateway Groups.

To view Gateway troubleshooting tools:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways > Gateway List**.

   The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*.
3. Select the required Gateway or Gateway Group.

   The *Gateways Overview* page appears for the selected Gateway/Gateway Group.
4. Click the **Troubleshooting** tab.

   The *Gateway Troubleshooting* page appears, showing troubleshooting tools and options for the selected Gateway/Gateway Group.

The page content is affected by your choice of task in the **Select Troubleshooting Task** drop-down menu:

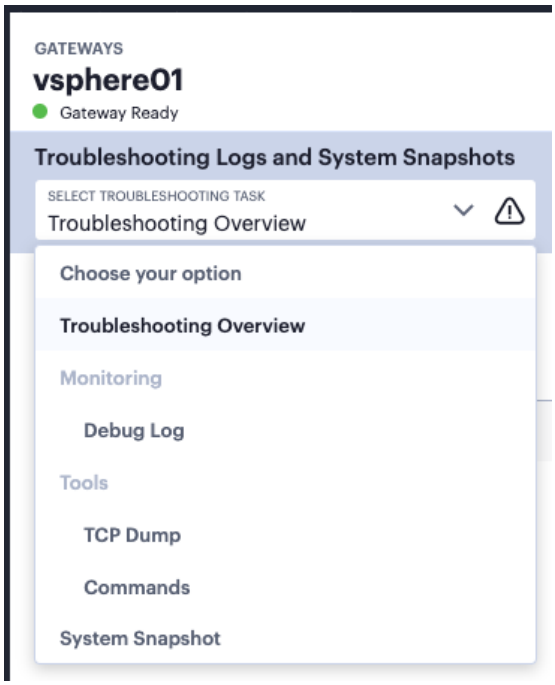

Viewing the list of available troubleshooting tasks for this Gateway

Choose from:

- **Troubleshooting Overview**: a list of all captured data based on your completed troubleshooting activities.
- **Debug Log**: make a trace recording of a selected process.
- **TCP Dump**: observe and record TCP packets on the network.
- **Commands**: run various common network troubleshooting commands.
- **System Snapshot**: take a system snapshot.

> The availability of specific tasks is limited by Gateway version. Gateway instances based on versions earlier than 21.6 might have access only to a subset of this list.

Troubleshooting Overview

The *Troubleshooting Overview* contains a list of all captured data based on your completed troubleshooting activities. This shows:

- TCP dumps
- Debug logs
- System snapshots

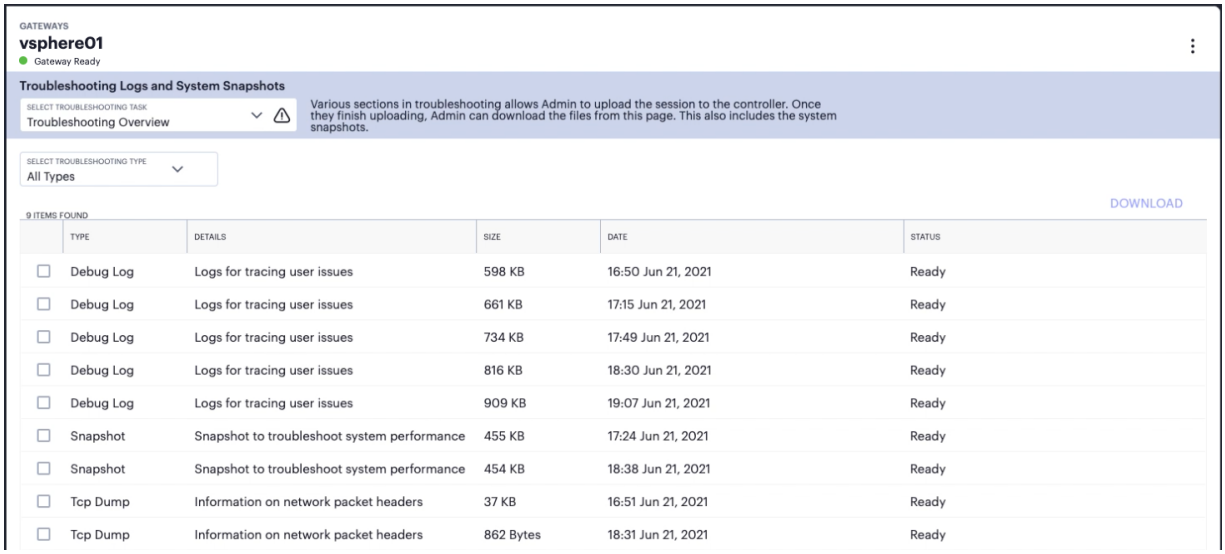

View a list of data files captured from previous troubleshooting activities

To filter the displayed files to a single type, use the **Select Troubleshooting Type** selector.

Debug Log

If your users are having difficulties getting access to a website or other resource, use the *Debug Log* task to make a trace recording for one or more selected processes running on the Gateway. This recording can then be uploaded to the *Controller* for offline analysis by *Ivanti* Technical Support.

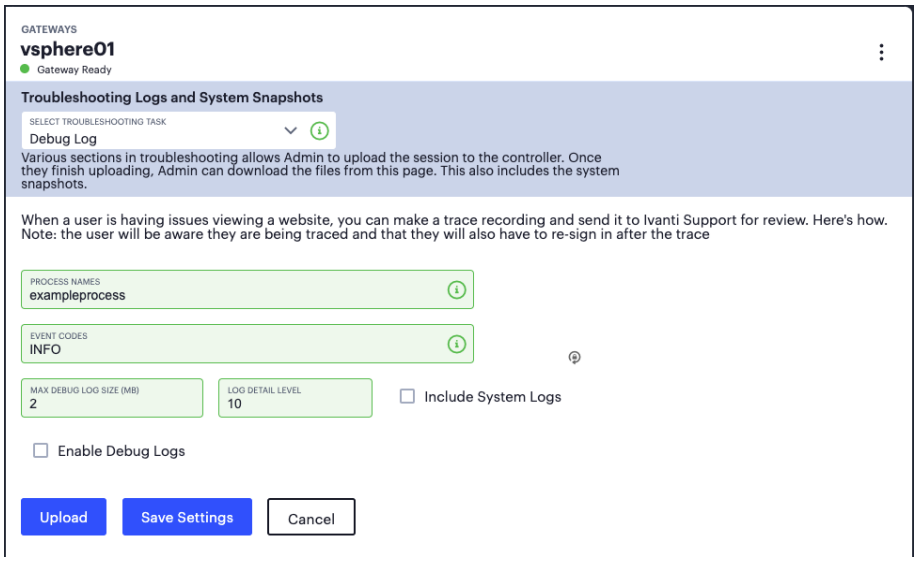

Recording a trace of a named process

Set the following fields:

- **Process Names**: the names of one or more Gateway processes to trace
- **Event Codes**: the event code, or codes, to observe in the trace. Only events matching these codes are appended to the log.
- **Max Debug Log Size (MB)**: The maximum allowable size, in MB, for the log file. Choose a value between 0-250.
- **Log Detail Level**: The level of verbosity to include in the log. Choose a value between 0-60, from lowest to highest detail.
- **Include System Logs**: Whether or not to include Gateway system logs (event, access, and admin logs) as part of the debug log.

To activate the trace, Select **Enable Debug Logs** and select **Save Settings**. To activate the trace, select **Enable Debug Logs** and then select **Save Settings**.

Use **Cancel** to return the settings on this page to the state last saved.

After a trace has been captured, select **Upload** to upload the trace log file to the *Controller*.

TCP Dump

Use the *TCP Dump* task to inspect packet data passing through the network interfaces on your Gateway:

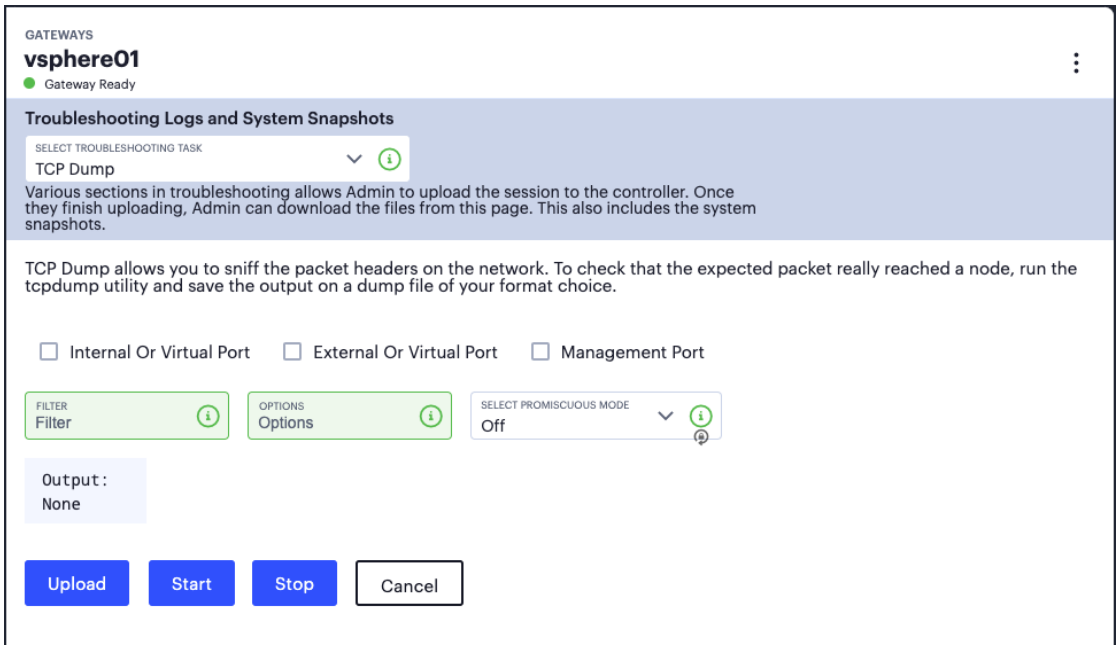

Inspecting TCP data packets

Select the network interfaces you want to inspect:

- Internal or Virtual Port
- External or Virtual Port
- Management Port (where applicable)

Next, set the following optional fields as required:

- **Filter**: A comma-separated filter expression for the TCP dump, based on standard UNIX TCP Dump filters. For more information, see https://www.freebsd.org/cgi/man.cgi?query=tcpdump.

  The following table provides some examples:

Example TCP dump filter expressions

Expression	Result
tcp port 80	Sniffs packets on TCP port 80

Expression	Result
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

- **Options**: Command-line options to specify with the tcpdump. Use a space-separated list in the form "-<option>". For more information on possible options, see https://www.freebsd.org/cgi/man.cgi?query=tcpdump.

  The following options are NOT supported: "-C", "-B", "-D", "-G", "-F", "-V", "-w", "-W", "-r", "-E", "-h", "-I", "-J", "-L", "-m", "-p", "-U", "-z", "-Z", "-O", NULL.

- **Promiscuous Mode**: Enable or disable promiscuous mode for the data capture.

To start the TCP dump, select **Start**. Then, select **Stop** to end the process when required.

The current state is reflected in the *Output* window.

Use **Cancel** to ret-urn the settings on this page to the state last saved.

After a TCP dump has been captured, select **Upload** to upload the log file to the *Controller*.

Commands

Use the *Commands* task to trigger common networking troubleshooting commands from the Gateway. For example, using the ping command:

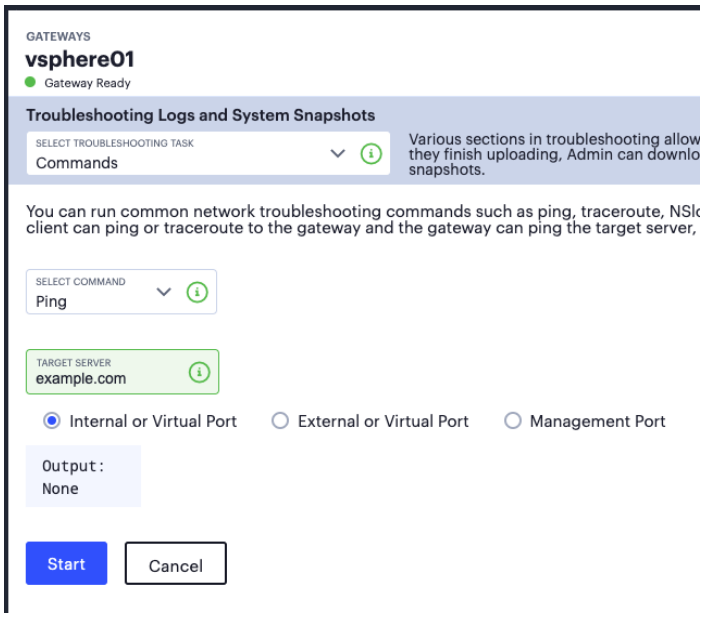

Running the Ping command

Use **Select Command** to select the command you want to run. Then, specify the parameters you want to use with that command.

To run the selected command with the specified parameters, select **Start**.

The command output is displayed in the *Output* window.

> Use **Cancel** to return the settings on this page to the state last saved.

The following table lists the available commands with parameters applicable in each case:

Commands

Command	Parameter	Description
ping	Target Server	The target hostname or IP address
	Gateway Interface	The network interface through which to capture the command output: internal, external, or management (where applicable)

Command	Parameter	Description
nslookup	Query Type	The *nslookup* query type you want to use: ANY, A, PTR, CNAME, MX, NS, SOA, TXT, UINFO, WKS
	Query	The target IP address or FQDN.
ARP	Target Server	The target hostname or IP address
	Gateway Interface	The network interface through which to capture the command output: internal, external, or management (where applicable)
Trace route	Target Server	The target hostname or IP address
	Gateway Interface	The network interface through which to capture the command output: internal, external, or management (where applicable)
Portprobe	Target Server	The target hostname or IP address
	Target Port	Specify a valid port at the target server in the range 1-65535.
	Select Protocol	Choose a protocol to use: TCP or UDP.
	Probe Count	The number of probes to send. Use a value between 1-100.
	Probe Timeout	The timeout value, in seconds, for each probe. Use a value between 1-180.

System Snapshot

Use the *System Snapshot* task to trigger a snapshot of your Gateway system performance as it is currently configured:

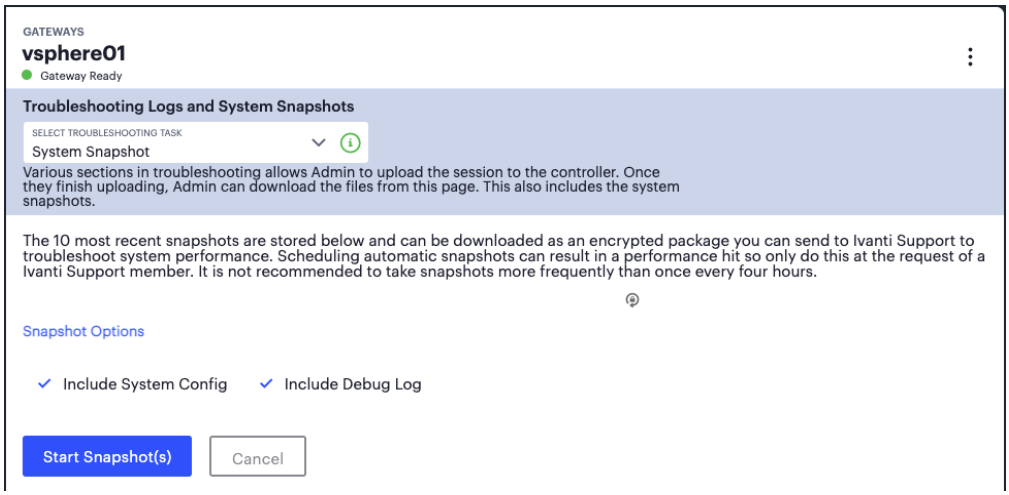

Taking a system snapshot

> Regular snapshots can result in a performance hit. *Ivanti* recommends not taking snapshots more frequently than once every four hours, unless at the request of a support representative.

Set the following fields:

- **Include System Config**: Choose whether or not to include the system config with the snapshot. This includes information such as host name, cache entries on the gateway, the *Controller* it is connected to, and so on.
- **Include Debug Log**: Choose whether or not to include the current debug log file with the snapshot.

> The current debug log is either created specifically by the admin through the *Debug Log* task, or in the absence of a created log, the default system debug log (generated automatically with Log level 0).

To capture the system snapshot, select **Start Snapshot(s)**.

> Use **Cancel** to return the settings on this page to the state last saved.

Completed snapshots are automatically uploaded to the *Troubleshooting Overview* task page and can be downloaded as an encrypted package to send to *Ivanti* Technical Support to troubleshoot system performance.

Adding Gateway Groups for High Availability

nZTA uses **Gateway Groups** to implement *high availability* for your Gateway deployments. Grouping Gateways together ensures the *Controller* can synchronize out policy changes to all Gateways in a group, automatically. To learn more about high availability and how it can benefit your application delivery, see Using Gateway Groups for High Availability.

To learn more about the process of adding and registering individual Gateways with the *Controller*, see the following workflows:

- To add a Gateway in VMware vSphere, see Workflow: Creating a Gateway in VMware vSphere.
- To add a Gateway in Amazon Web Services (AWS), see Workflow: Creating a Gateway in Amazon Web Services.
- To add a Gateway in Microsoft Azure, see Workflow: Creating a Gateway in Microsoft Azure.
- To add a Gateway in KVM/OpenStack, see Workflow: Creating a Gateway in KVM/OpenStack.
- To add a Gateway in Google Cloud Platform (GCP), see Workflow: Creating a Gateway in Google Cloud Platform.

After you have created your Gateways, you can add them to a Gateway Group. A Gateway already added a group cannot be added to a further Gateway Group, and cannot be used by more than one Secure Access Policy.

To add a Gateway Group:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways**.

   The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*.
3. To add a new Gateway Group, select **Create** from the top-right:

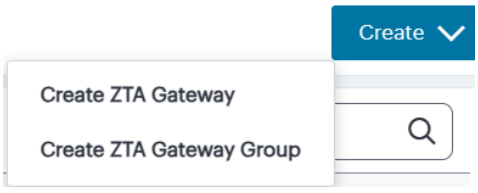

Add a new Gateway or Gateway Group

4. In the drop-down menu, click **Create ZTA Gateway Group**.

   The **Create ZTA Gateway Group** dialog appears.

5. Enter a **Name** for the Gateway Group.

6. (Optional) Enter a **Description** for the Gateway Group.

7. Enter the **Load Balancer IP addresses or CNAME**. That is, a list of the public IP addresses or CNAMEs of the load balancer's front-end interface that your end users connect to. Enter a value, then select **Add** to add it to the list. Repeat this step for each entry you want to add.

8. (Optional) Select the pre-existing Gateways you want to add to this group. Use the drop-down list to select a Gateway, then select **Add**. Repeat this step for each Gateway you want to add.

9. To add this Gateway Group, click **Create Gateway Group**.

A newly added Gateway Group appears on the *All Gateways* page alongside all *ungrouped* Gateways.

> **For cloud-based Gateways**: If you plan to add multiple deployed Gateways to a Gateway Group for high availability, additional configuration is required for the public IP addresses assigned to your instances. For more details, see Registering an Amazon Web Services Gateway (for AWS) or Creating a Gateway using the Azure Template and Image Files (for Azure).

Creating Gateway Selectors

The Gateway Selector feature provides optimal gateway selection and dynamic failover when deploying multiple geographically located nZTA gateways.

Tenant admins can then manually configure policies that determine to which Gateway an end-user is sent when they access an application. Tenant admins can select Gateways/Gateway Groups and set the priorities to identify to which set of Gateways the Client should connect to access the application.

To add a Gateway Selector:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.

2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways > Gateway Selector**.

   The *Gateways Selector List* page appears.

3. To add a new Gateway Selector, click **Create Gateway Selector**.

   The *Gateway Selector Details* page appears.

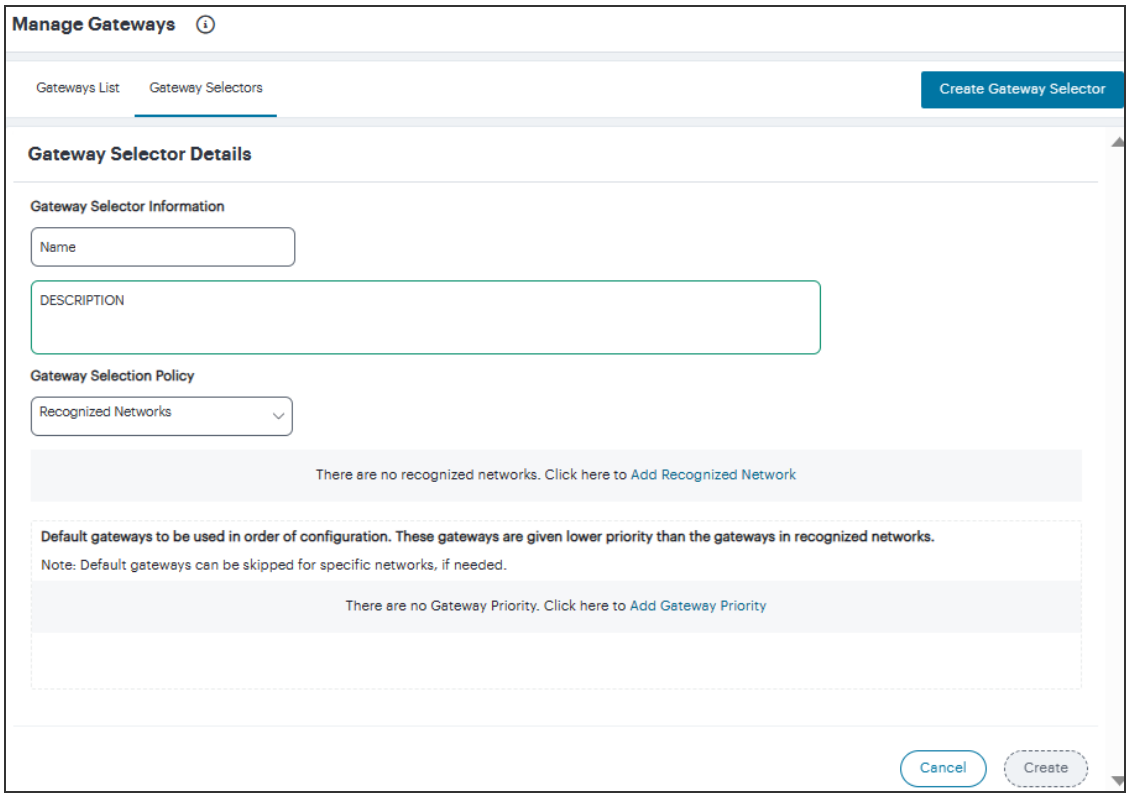

   Gateway Selector Details

4. Enter a **Name** for the Gateway Selector.
5. (Optional) Enter a **Description** for the Gateway Selector.
6. In the *Gateway Selector Policy* section, you can choose or add **Recognized Networks**, and choose or add gateway priority.

7. To add a Recognized Network:

    1. Click the **Add Recognized Network** link.

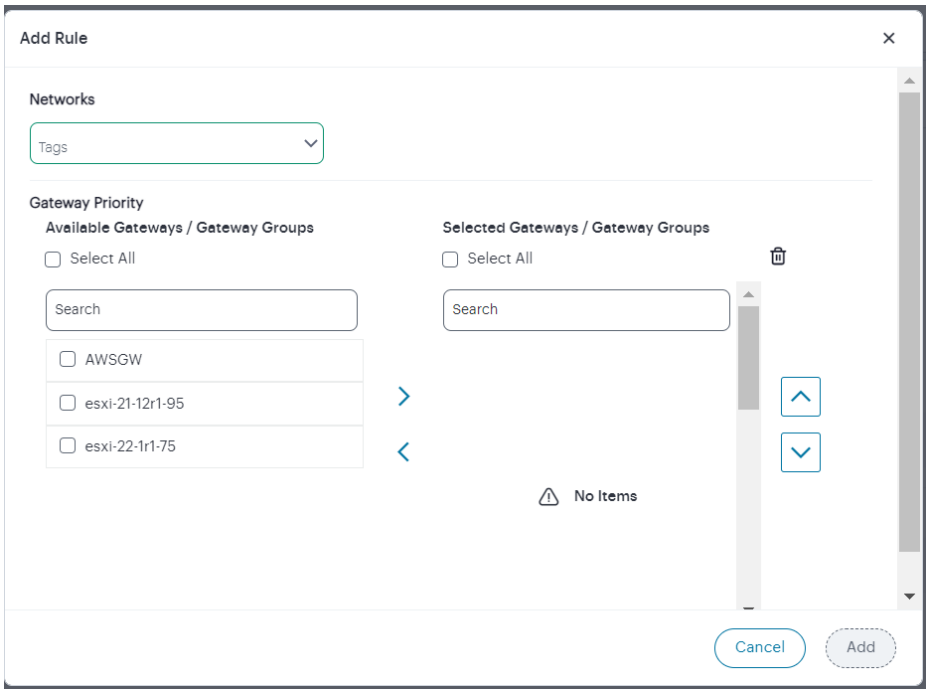

Add recognized network

    2. In the *Add Rule* dialog, from the **Networks** drop-down list select a tag for the recognized network. To learn about creating tags, see the *Associating Geographical locations to IP Addresses* section in Using the Insights Menu to Monitor User Activity and Service Usage.

    3. In the Gateway Priority section, select the gateways/gateway groups from the **Available Gateways / Gateway Groups** list.

    4. Use the Up/Down arrows to change the priorities.

    5. Click **Add**.

        A newly added Recognized Network appears in the Gateway Selection Policy section.

8. To add Gateway Priority:

    1. Click the **Add Gateway Priority** link.

       The Gateway Priority window appears.

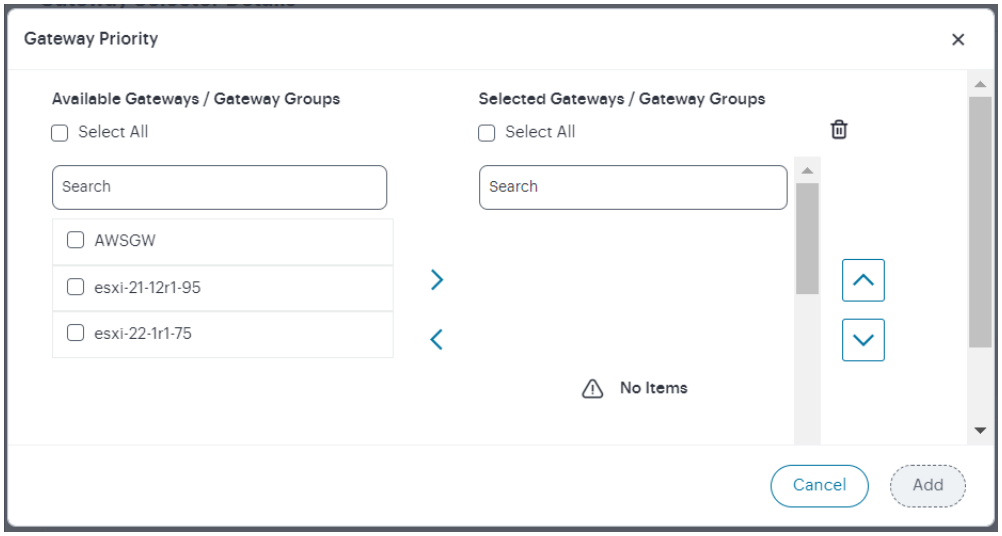

Set Gateway priority

    2. Select the gateways/gateway groups from the **Available Gateways / Gateway Groups** list.

    3. Use the Up/Down arrows to change the priorities.

    4. Click **Add**.

       A newly added Gateway Priority appears in the Gateway Selection Policy section.

9. In the Gateway Selector Details page, click **Create**.

   A newly added Gateway Selector appears in the Gateway Selectors page.

10. To modify a Gateway Selector, click the adjacent three dots, then select **Edit**. In the Gateway Selector Details page, make the necessary changes and click **Update**.

11. To remove a Gateway Selector, click the adjacent three dots, then select **Delete**, and then confirm by clicking **Delete**.

Workflow: Creating a Gateway in VMware vSphere

The process of registering a vSphere Gateway with *nZTA* involves two main procedures, to be completed in sequence:

- Create the Gateway record in the *Controller*.
- Create the Gateway virtual machine instance in VMware vSphere.

After these steps have been completed successfully, the *Controller* and Gateway establish communication with each other.

Before you start, make sure that you have the following information and files for the Gateway:

- An identifying name for the Gateway
- The public IP address for the Gateway. This is the IP address at which clients can externally reach the Gateway instance.
- The Gateway geographic location
- (Optional) The name of the *Gateway Group* to which you want to add this new Gateway record. To learn more about Gateway Groups, see Adding Gateway Groups for High Availability.
- The Gateway OVF template: https://pulsezta.blob.core.windows.net/gateway/ISA-V-VMWARE-ZTA-22.7R1-371.1.zip

  Download a copy of the OVF template archive file and unpack to a local workstation. Make sure the resulting file set is accessible from the vSphere Console.

  You can also choose to download this file from the **Gateways Overview** page in the *nZTA* Tenant Admin Portal. The opportunity to do this occurs later in this process.

- Credentials for the vSphere Console.

  These credentials must include sufficient permissions to create a VM from a template image.

  By default, *nZTA* derives Gateway DNS and network interface settings through DHCP (provided this is configured in your vSphere environment). If, instead, you want to manually specify your Gateway DNS and network interface settings, make sure you have the following additional information:

- The internal/private subnet IP address, subnet mask, and network gateway IP address.
- The primary (and optional secondary) DNS server IP address, and search domain.

- The external interface IP address, subnet mask, and network gateway IP address.
- (Optional) The management interface IP address, subnet mask, and network gateway IP address.

If you choose to deploy a vSphere-based Gateway instance with DHCP configuration, DNS server settings are not configurable through the *nZTA* Tenant Admin UI. In this scenario, make sure you have at least a primary DNS server configured and available through DHCP (a secondary DNS server is optional). Make sure also that a DNS search domain is properly configured.

Adding a VMware vSphere Gateway

To register a Gateway on your *Controller*, use the **Gateway Details** dialog.

To begin, log into the *Controller* Tenant Admin Portal using the credentials provided in your welcome email. Two outcomes are possible:

- On unconfigured *nZTA* systems, the *Secure Access Setup* Onboarding wizard appears (see Working with the Onboarding Wizard). In this case, click **Add Gateway**.
- On configured *nZTA* systems, the *Network Overview* page appears. In this case:

  1. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways**.

     The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*.

  2. To add a new Gateway, select **Create** from the top-right:

     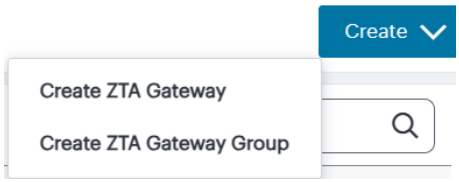

     Add a new Gateway or Gateway Group

  3. In the drop-down menu, click **Create ZTA Gateway**.

In both cases, the **Gateway Details** dialog appears.

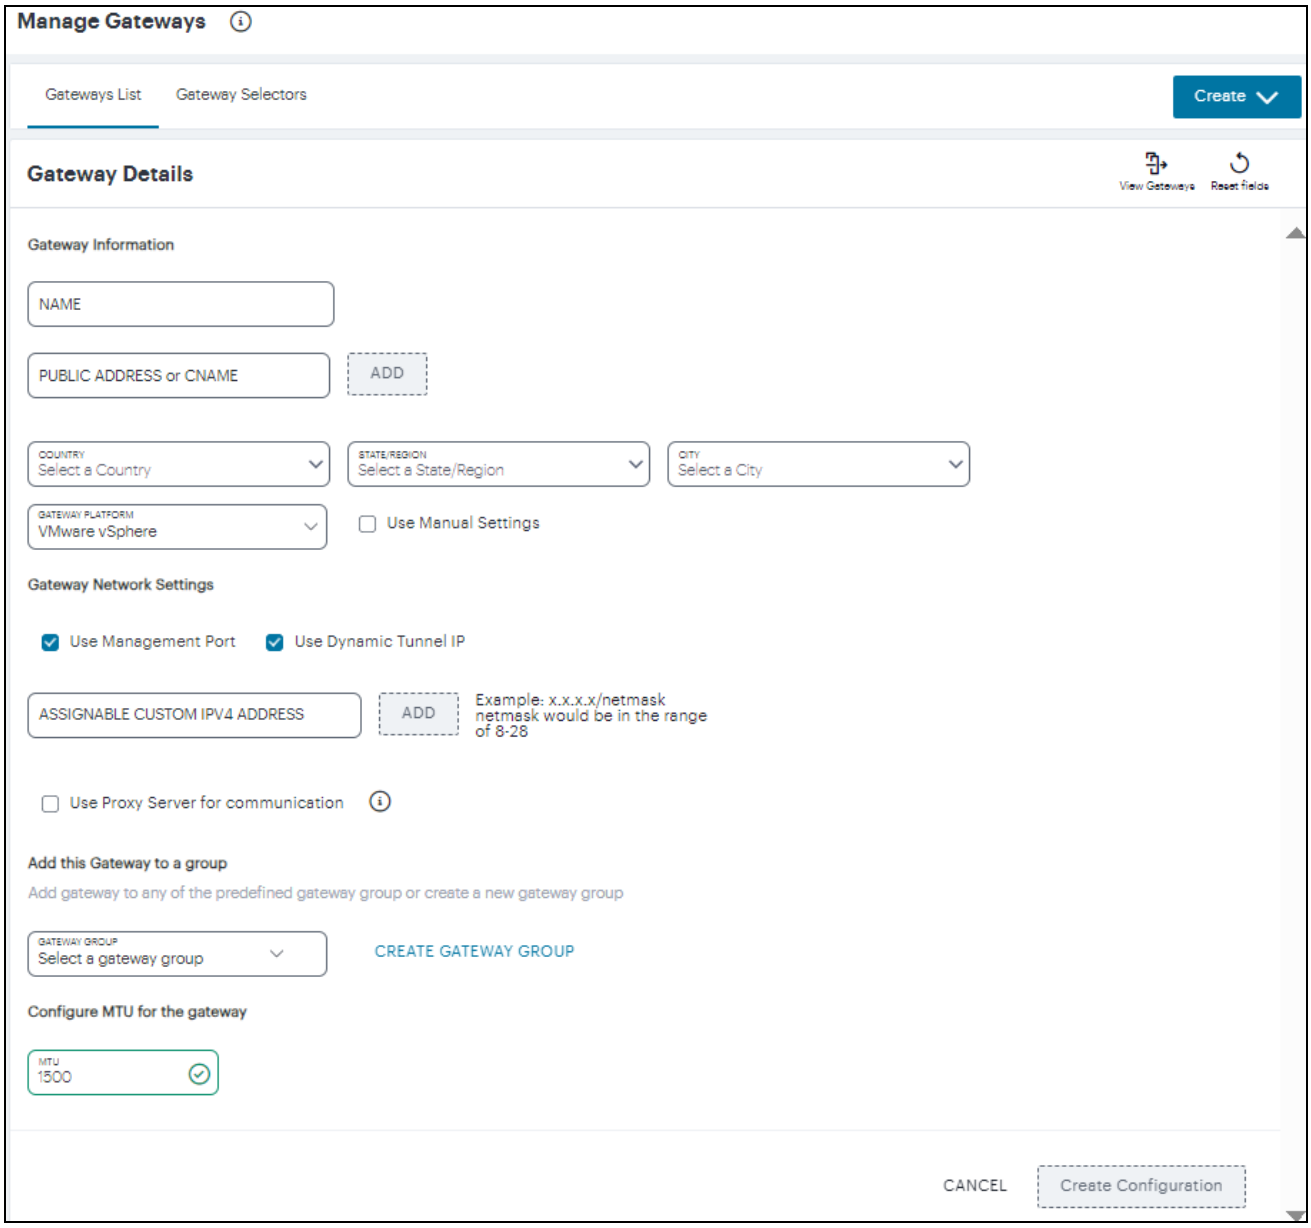

Gateway Details

Enter the following details:

1. Enter a **Name** for the Gateway.
2. Enter one or more **Public Address or CNAME** (Public IP address or CNAME) for the Gateway. Select **Add** to add each entry to the list. To learn more about this setting, see Configuring Networks in your Gateway Datacenter.
3. Select the geographic location details for the Gateway.

4. For **Gateway Platform**, select "VMware vSphere".
5. (Optional) To enter your vSphere Gateway instance DNS and network interface settings manually, select **Use Manual Settings**. To instead allow *nZTA* to use DHCP-derived settings for DNS and network interfaces, leave **Use Manual Settings** un-selected. To learn more about these settings, see Configuring Networks in your Gateway Datacenter.
6. (Optional) Select the **Use Management Port** check box to use management network ports for *nZTA* traffic rather than internal ports.

   When the management port is enabled, Gateway will use management interface to communicate with *Controller* and NTP Server.
   The Gateway will still use the internal port for DNS resolution and NTP server name resolution.
   If the internal DNS cannot resolve the *Controller* domain, the internal interface will require internet access.

7. (Optional) Select the **Use Dynamic Tunnel IP** check box to configure a pool of IP addresses that are dynamically mapped to client sessions with this Gateway, such that user traffic from the Gateway to an application can be identified as originating from a specific client. To learn more, see Using Dynamic IP Addressing to Profile Client Traffic.

   The *Custom IP Pool* dialog appears:

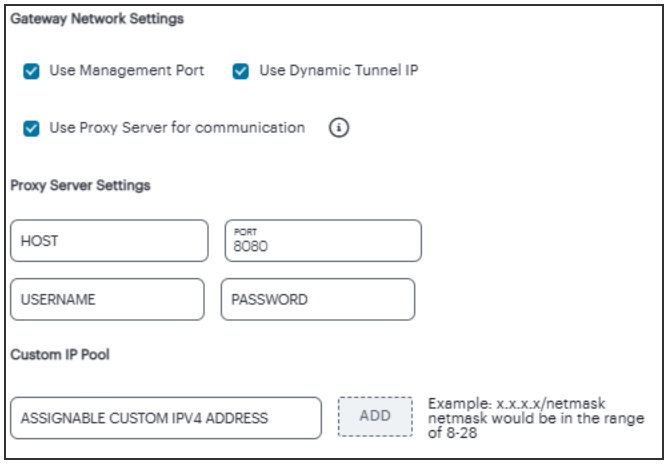

   Gateway Details - Custom IP Pool settings

   Dynamic Tunnel IP addresses are not supported in Gateway Groups.

Use the **Assignable Custom IPv4 Address** field to enter an IP address and subnet (in the range 8-28) in CIDR notation, then click **Add**. Repeat this step for each address/subnet you want to use.

8. (Optional) Select the **Use proxy server for communication** check box to enable nZTA to *Controller* communication via proxy server. Proxy is supported on both internal and management interfaces of the Gateway. Once enabled, enter host name and port. Optionally, if your proxy server requires further authentication, enter a username and password to log in to the proxy server.

   

   Admin can configure proxy for existing Gateway after upgrading it to 22.4R3 version or later.

9. (Optional) Select a **Gateway Group** to which the new Gateway is to be added. To learn more about Gateway Groups, see Adding Gateway Groups for High Availability.

10. If you elected to use manual settings, the following panel appears:

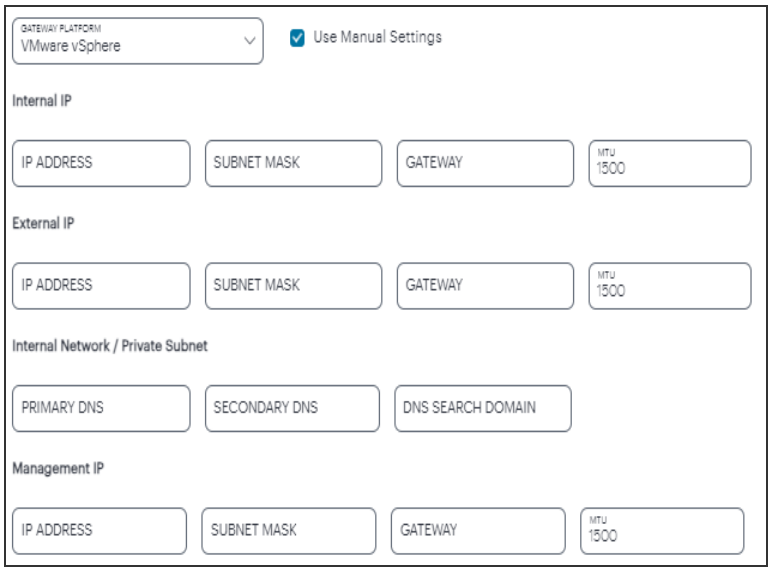

Gateway Network Configuration - manual settings

Enter the following details:

1. Specify the internal **IP Address** for the Gateway.
2. Specify the internal **Subnet Mask** for the Gateway.
3. Specify the internal network gateway IP address as the **Gateway** setting.
4. Specify the **MTU** size between 576 and 1500.
5. Enter the **Primary DNS** IP address for the Gateway.
6. (Optional) Enter the **Secondary DNS** IP address for the Gateway.
7. Enter the **DNS Search Domain** for the Gateway.
8. Specify the external **IP Address** for the Gateway.
9. Specify the external **Subnet Mask** for the Gateway.
10. Specify the external network gateway IP address as the **Gateway** setting.
11. Specify the **MTU** size between 576 and 1500.

12. Specify the management **IP Address** for the Gateway.

    Management network settings are optional, unless the **Use Management Port** check box is selected.

13. Specify the management **Subnet Mask** for the Gateway.
14. Specify the management network gateway IP address as the **Gateway** setting.
15. Specify the **MTU** size between 576 and 1500.

11. (Optional) Select a **Gateway Group** to which the new Gateway is to be added. To learn more about Gateway Groups, see Adding Gateway Groups for High Availability.
12. **Configure MTU for the gateway**: Configurable MTU size allows admin to modify the default setting of nZTA gateways wherever it is needed. The value allowed is in the range of 576 to 1500 (IPv4).
13. To add a Gateway definition based on the settings you specified in this dialog, select **Create Configuration**.

After you complete this process, an unregistered Gateway record is created on the *Controller*. You can view this Gateway record on the **Manage Gateways > Gateways List** page.

Next, create the Gateway virtual machine instance in VMware vSphere and allow it to register with the *Controller*. This links the Gateway record on the *Controller* with the actual Gateway virtual machine, see Registering a VMware vSphere Gateway.

Registering a VMware vSphere Gateway

This section describes the steps necessary to create the Gateway virtual machine in the vSphere console. To learn more about the operations included here, refer to VMware's own documentation for full details.

For reference, the recommended minimum requirements for a Gateway virtual machine instance in vSphere are:

- 4 vCPU's and 8 GB memory, or
- 8 vCPU's and 32 GB memory

After you have created an unregistered Gateway record on the **Gateways** page, you must create the Gateway instance on the *VMware vSphere Console*. This process facilitates communication between the *Controller* and the Gateway instance.

Before you start, make sure you have obtained the Gateway definition file from the *Controller*. This definition file includes the settings necessary to configure the new Gateway virtual machine with the identity and location of the *Controller*, and is used during the registration process described later in this section.

To download the Gateway definition file:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways**.

   The *Gateways List* page appears.
3. Locate and select your unregistered vSphere Gateway record from the list of available Gateways.
4. Click the **Download** icon and select **Download gateway init config** to obtain a copy of the Gateway definition file.

   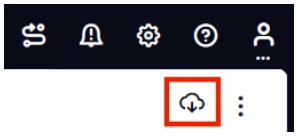

   The Download Icon
5. Save the downloaded text file to a location accessible from the vSphere Console.

   > The Gateway definition file is valid for 24 hours. If this period expires, you must replace the Gateway to generate a new Gateway definition file.
6. (Optional) If you have not yet downloaded the latest version of your Gateway VM file, click the **Download** icon and select **Download gateway VM image**. Save the archive file and unpack to a local workstation. Make sure the resulting file set is accessible from the vSphere Console.

To register a Gateway:

1. Access the vSphere console, either from a client or a web browser.
2. In the vSphere console, start the *Deploy OVF Template* wizard to create a new virtual machine based on the *nZTA* vSphere Gateway template.

3. In the wizard:
    - Choose to deploy from a local file.
    - Locate and upload your OVF/VMDK template files.
    - Provide an identifying name and location for the new Gateway virtual machine.
    - Choose any required compute resource.
    - Choose the required storage settings.
    - Customize the *vApp properties* of your virtual machine and, in the **VA IVE Configuration** parameter, paste the raw text of your Gateway definition file. To obtain the Gateway definition file, see the process described earlier in this section.
    - Confirm all settings.
    - Finish the wizard to create the Gateway VM.
4. Locate the new Gateway VM in the hosts and clusters.
5. Start the Gateway VM by powering it on.
6. Wait until the power-up is complete.
7. Return to the **Gateways List** page on the *Controller*.
8. Locate the new Gateway record in the list and confirm that its status has updated to *Connected*.
9. (Optional) After you have registered a Gateway, you can configure it (or the Gateway Group to which it belongs) as the default Gateway if required. A default Gateway handles all requests from applications that are not referenced by any secure access policy. See Configuring a Default Gateway for Application Discovery.

    After this process completes, you can move to the next stage of *nZTA* configuration, which is Working with User Authentication.

    In case of registration failure due to Gateway configuration mistakes in firewall rules, DNS, etc., you can re-register the gateway. It does not require re-deploying of Gateway. For details see "Re-registering an Amazon Web Services Gateway" on page 264

Re-registering a VMware vSphere Gateway

Re-registration of Gateway includes the following capabilities:

- Gateway triggers re-registration on every launch of the gateway if the registration with controller fails or any update is made to the configuration parameters.
- The "View registration error report" option provides the reason for failure and solution to rectify it.

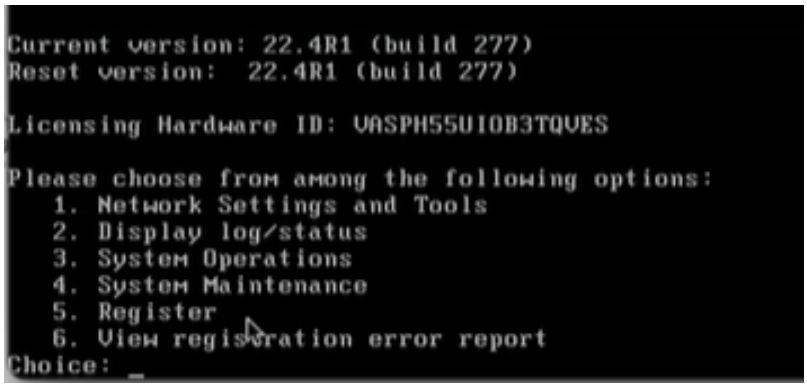

- On registration failures, admin is provided with the "Register" option to trigger the registration manually along with the existing debugging options such as networking tools, reboot, etc. This option can be used after rectifying any external issues such as network reach issue or Firewall rules following controller traffic from Gateway.
- To rectify registration failure due to the config error, first update the config settings in the Controller and download the config file. Then shut down Gateway VM, update downloaded config, and then boot Gateway.

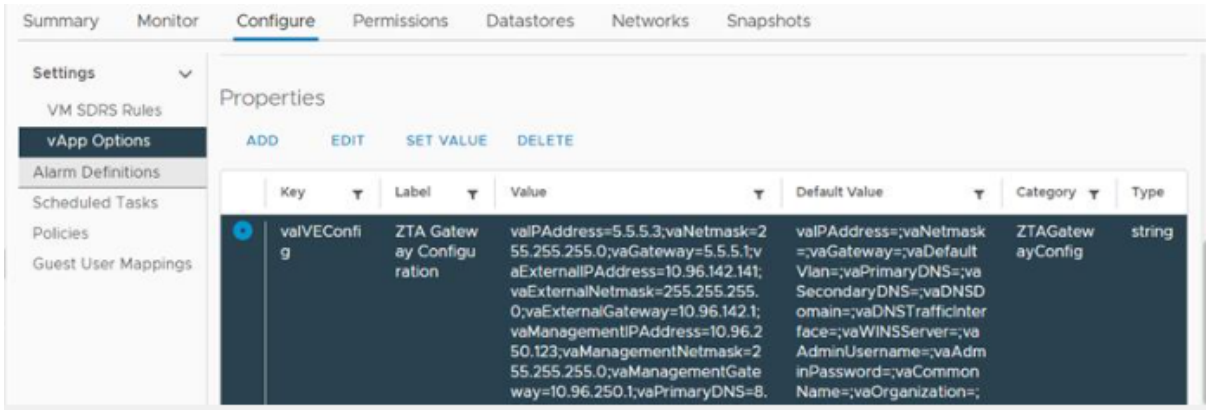

Sample screen: Update Config Value

An option is provided to regenerate and download the gateway init config from the controller admin interface.

Workflow: Creating a Gateway in Amazon Web Services

The process of registering an AWS Gateway with *nZTA* involves two main procedures, to be completed in sequence:

- Create the Gateway record in the *Controller*.
- Create the Gateway virtual machine instance in Amazon Web Services (AWS).9

After these steps have been completed successfully, the *Controller* and Gateway establish communication with each other.

Before you start, make sure that you have the following information and files for the Gateway:

- An identifying name for the Gateway
- The public IP address for the Gateway. This is the IP address at which clients can externally reach the Gateway instance, typically an elastic IP address provided by AWS.
- The Gateway geographic location.
- (Optional) The name of the *Gateway Group* to which you want to add this new Gateway record. To learn more about Gateway Groups, see Adding Gateway Groups for High Availability.
- The primary (and optional secondary) DNS server IP address, and search domain.

- The Gateway template file. *nZTA Gateway*s can be deployed in a new VPC or an existing VPC, using the **Nitro** hypervisor. Select the JSON template file that is applicable to your requirements:
    - To deploy in an existing VPC - Nitro hypervisor (M5-type instances):

      https://pulsezta.blob.core.windows.net/gateway/templates/AWS/24-1-371/ivanti-2nic-existing-vpc.json

      https://pulsezta.blob.core.windows.net/gateway/templates/AWS/24-1-371/ivanti-3nic-existing-vpc.json

    - To deploy in a new VPC - Nitro hypervisor (M5-type instances):

      https://pulsezta.blob.core.windows.net/gateway/templates/AWS/24-1-371/ivanti-2nic-new-vpc.json

      https://pulsezta.blob.core.windows.net/gateway/templates/AWS/24-1-371/ivanti-3nic-new-vpc.json

If you want to use a Management interface, you must download and use the 3 NIC template.

You might not be able to specify the download location given here directly to AWS. In this case, download the Gateway template file first to your local workstation and specify this location instead.

You can also choose to download this file from the **Gateways Overview** page of the *nZTA* Tenant Admin Portal. The opportunity to do this occurs later in this process.

- The Gateway AMI identifier. *nZTA* gateway AMIs are available in all AWS regions (except China).

  In most cases, after the Gateway image becomes available in the AWS Marketplace, the AMI applicable to your region is automatically selected. Alternatively, to perform a manual search, follow these steps:

  1. Log into the AWS console.
  2. Navigate to **EC2 > Images > AMIs**.
  3. Select "Public Images".
  4. Search for the image corresponding to your selected hypervisor:
     - Nitro: "ISA-V-NITRO-ZTA-22.7R1-371.1.img"
  5. Make a note of the corresponding AMI ID.

- Credentials for the AWS Management Console.

  

  These credentials must include sufficient permissions to create a stack.

- The SSH key pair file that you are using with the AWS Management Console.

Adding an Amazon Web Services Gateway

To registering a Gateway on your *Controller*, use the **Gateway Details** dialog.

To begin, log into the *Controller* Tenant Admin Portal using the credentials provided in your welcome email. Two outcomes are possible:

- On unconfigured *nZTA* systems, the *Secure Access Setup* Onboarding wizard appears (see Working with the Onboarding Wizard). In this case, click **Add Gateway**.

- On configured *nZTA* systems, the *Network Overview* page appears. In this case:

    1. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways**.

       The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*.

    2. To add a new Gateway, select **Create** from the top-right:

       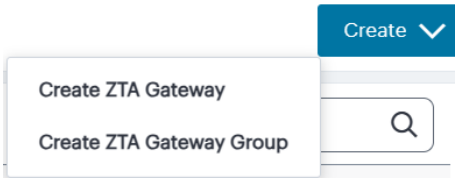

       Add a new Gateway or Gateway Group

    3. In the drop-down menu, click **Create ZTA Gateway**.

In both cases, the **Gateway Details** dialog appears.

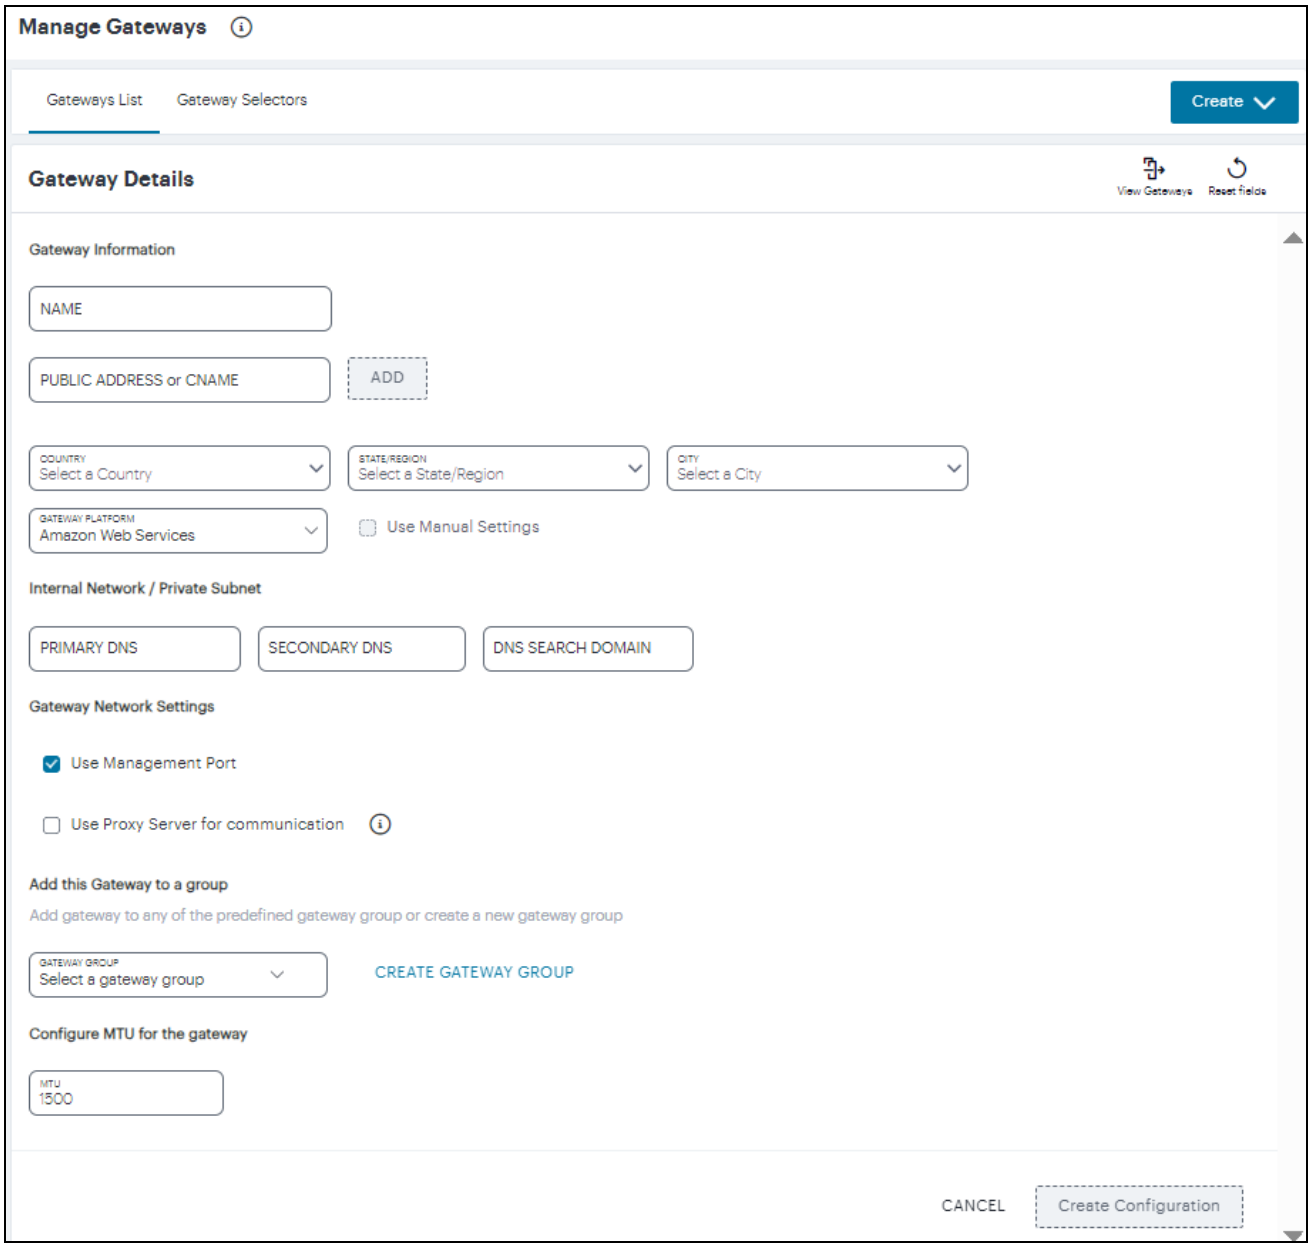

Gateway Details

Enter the following details:

1. Enter a **Name** for the Gateway.
2. Enter one or more **Public Address or CNAME** (Public IP address or CNAME) for the Gateway. Select **Add** to add each entry to the list. To learn more about this setting, see Configuring Networks in your Gateway Datacenter.

3. Select the geographic location details for the Gateway.
4. For **Gateway Platform**, select "Amazon Web Services".
5. Enter the Primary DNS IP address for the Gateway.
6. (Optional) Enter the Secondary DNS IP address for the Gateway.
7. Enter the DNS Search Domain for the Gateway.

   Make sure the specified DNS service can resolve the IP address of your *Controller*. Issues here can cause registration of the Gateway with the *Controller* to fail.

8. (Optional) Select the **Use Management Port** check box to use management network ports for *nZTA* traffic rather than internal ports.

   When the management port is enabled, Gateway will use management interface to communicate with *Controller* and NTP Server.
   The Gateway will still use the internal port for DNS resolution and NTP server name resolution.
   If the internal DNS cannot resolve the *Controller* domain, the internal interface will require internet access.

9. (Optional) Select the **Use proxy server for communication** check box to enable nZTA to *Controller* communication via proxy server. Proxy is supported on both internal and management interfaces of the Gateway. Once enabled, enter host name and port. Optionally, if your proxy server requires further authentication, enter a username and password to log in to the proxy server.

   Admin can configure proxy for existing Gateway after upgrading it to 22.4R3 version or later.

10. (Optional) Select a **Gateway Group** to which the new Gateway is to be added. To learn more about Gateway Groups, see Adding Gateway Groups for High Availability.
11. **Configure MTU for the gateway**: Configurable MTU size allows admin to modify the default setting of nZTA gateways wherever it is needed. The value allowed is in the range of 576 to 1500 (IPv4).

To add a Gateway definition based on the settings you specified in this dialog, select **Create Configuration**.

After you complete the first part of this workflow, an unregistered Gateway record is created on the *Controller*. This Gateway record can be seen on the **Manage Gateways > Gateways List** page.

Next, create the Gateway virtual machine instance on AWS and allow it to register with the *Controller*. This links the Gateway record on the *Controller* with the actual Gateway virtual machine, see Registering an Amazon Web Services Gateway.

Registering an Amazon Web Services Gateway

This section describes the steps necessary to create the Gateway virtual machine in the AWS management console. To learn more about the operations included here, refer to Amazon's own documentation for full details.

After you have created an unregistered Gateway record on the **Gateways** page, you must create the Gateway instance on the *AWS Management Console*. This process facilitates communication between the *Controller* and the Gateway instance.

Before you start, make sure you have obtained the Gateway definition file from the *Controller*. This definition file includes the settings necessary to configure the new Gateway virtual machine with the identity and location of the *Controller*.

To download the Gateway definition file:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways**.

   The *Gateways List* page appears.
3. Locate and select your unregistered AWS Gateway record from the list of available Gateways.
4. Click the **Download** icon to obtain a copy of the Gateway definition file.

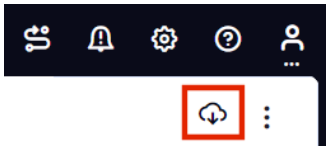

The Download Icon

5. Save the downloaded text file to a location accessible from the AWS Management Console.

   > The Gateway definition file is valid for 24 hours. If this period expires, you must replace the Gateway to generate a new Gateway definition file.

6. (Optional) If you have not yet downloaded the latest version of your Gateway VM file, click the **Download** icon and select **Download gateway VM image**. Save the archive file and unpack to a local workstation. Make sure the resulting file set is accessible from the AWS Management Console.

To register a Gateway:

1. Access the AWS Management Console and log in using your credentials.
2. In the AWS **Services** menu, select **CloudFormation**.

   The **CloudFormation** home page appears.
3. Click **Create Stack** and then, from the sub-menu, select **With new resources (standard)**.

   The **Specify template** step of the **Create Stack** wizard appears.
4. Under **Prerequisite - prepare template**, select the **Template is ready** option.
5. Under **Specify Template**, select the **Upload a template file** template source option.
6. Under **Upload a template file**, click **Choose File** and select the Gateway template file that you downloaded at the start of this process.

   The file uploads, and the AWS S3 URL for the uploaded template file appears automatically.
7. Click **Next**.

   The **Specify stack details** step of the **Create Stack** wizard appears. This page displays the details and parameters required by the Gateway template.
8. Enter a **Stack name**.

9. Specify the parameters as appropriate for your deployment:
    - If you are deploying the Gateway instance into a *new VPC*, you can accept the default values used for all parameters.
    - If you are deploying the instance into an *existing VPC*, you must manually specify the details of your existing VPC into the parameters on the page. For more information, contact *Ivanti* Technical Support.
10. Under **ZTAGateway Configuration**, identify the Gateway AMI using its **ZTAGateway AMI ID**. Choose the designated AMI for the region in which you are deploying the Gateway instance.
11. For **Instance Type**, select the instance type that fits your hypervisor choice (Xen or Nitro) and minimum requirements, based on the following recommended types:

    For reference, the recommended minimum requirements for a Gateway virtual machine instance in AWS are:
    For Nitro hypervisor-based instances, use M5 types:

    - m5.large (2 vCPU, 8 GB Memory) (2NIC min)
    - m5.xlarge (4 vCPU, 16 GB Memory) (3NIC min)
    - m5.2xlarge (8 vCPU, 32 GB Memory)
    - m5.4xlarge (16 vCPU, 64 GB Memory)
12. For **ZTAGateway Config Data**, paste in the raw text of your Gateway definition file. To obtain the Gateway definition file, see the process described earlier in this section.
13. For **SSH Key Name**, specify your existing SSH key pair name.

14. For **Load Balancer Configuration**, If you plan to deploy multiple Gateways inside a Gateway Group, select "Yes" to deploy a new internet-facing network load balancer instance alongside the Gateway. Select "No" to launch only this Gateway instance.

    This option is applicable only for new VPC templates.

    If you elect to launch a load balancer, the following pre-configuration is applied:

    - An Elastic IP address is assigned to the load balancer.
    - A TCP listener is configured on port 443.
    - An IP-based Target Group is created and the private IP address of the deployed Gateway's external network interface is added as a target.
    - A health-check is configured on TCP port 443.
    - Stickiness is enabled on the Target Group.

    After you have deployed the Gateway and Load Balancer, you must return to the Tenant Admin Portal on the *Controller* and update the Gateway Group **Load Balancer IP ADDRESS** setting to be the Load Balancer's public IP address.

    If you want to configure the Load Balancer to balance across further Gateway instances from the Gateway Group, you must deploy each subsequent Gateway into an existing VPC and then update the Load Balancer Target Group.

    With new VPC templates, a NAT gateway is deployed for routing outbound Internet traffic from the Gateway's internal network interface in order for the Gateway to be able to reach the *Controller*.

    Public IP addresses are not automatically assigned to any of your Gateway's network interfaces. If you are deploying a Gateway into an existing VPC, in order for the Gateway to be able to reach the *Controller* from it's internal network interface, make sure you allow outbound Internet traffic from the Private Subnet for the deployed Gateway.

    To learn more about high availability and Gateway Groups, see Adding Gateway Groups for High Availability and Using Gateway Groups for High Availability.

15. Click **Next**.

    The **Configure stack options** step of the **Create Stack** wizard appears. All properties that were specified either in the template or in earlier steps are populated automatically.

    No changes or new inputs are required.

16. Click **Next**.

17. The **Review** step of the **Create Stack** wizard appears.

18. Confirm all displayed details.

19. Click **Create stack**.

    The **Stacks** page appears. The new stack is listed using the **Stack name** you specified during the wizard. The new stack has a status of *CREATE_IN_PROGRESS*.

20. Wait for the status of the new stack to reach *CREATE_COMPLETE*.

21. *(This step is required only if you have not deployed your Gateway with a Load Balancer or NAT at the front-end)* Elastic IP addresses are not automatically assigned to any of the Gateway's network interfaces. Therefore, before you can access the new Gateway instance from the *Controller*, you must associate a new Public IP address with the external interface of the Gateway. Then, return to the *Controller* Tenant Admin Portal and update the Gateway *Public IP Address* setting to match this address.

22. In the Tenant Admin Portal **Secure Access > Gateways > Gateways List** page, make sure the new Gateway has a confirmed status of *Connected*.

    

    You can directly access your AWS instance over SSH using *AWS EC2 Instance Connect*. To configure *AWS EC2 Instance Connect*, refer to the *Amazon Web Service Documentation*. You can then connect to the instance directly as a serial console using SSH from inside the *AWS Management Console*, refer to the *Amazon Web Service Documentation*.

23. (Optional) After you have registered a Gateway, you can configure it (or the Gateway Group to which it belongs) as the default Gateway if required. A default Gateway handles all requests from applications that are not referenced by any secure access policy. See Configuring a Default Gateway for Application Discovery.

After this process completes, you can move to the next stage of *nZTA* configuration, which is Working with User Authentication.

In case of registration failure due to Gateway configuration mistakes in firewall rules, DNS, etc., you can re-register the gateway. It does not require re-deploying of Gateway. For details see "Re-registering an Amazon Web Services Gateway" below.

Re-registering an Amazon Web Services Gateway

Re-registration of Gateway includes the following capabilities:

- Gateway triggers re-registration on every launch of the gateway if the registration with controller fails or any update is made to the configuration parameters.
- The "View registration error report" option provides the reason for failure and solution to rectify it.

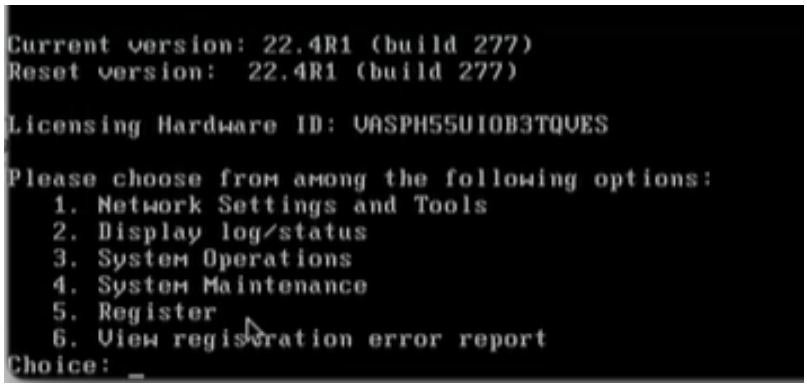

- On registration failures, admin is provided with the "Register" option to trigger the registration manually along with the existing debugging options such as networking tools, reboot, etc. This option can be used after rectifying any external issues such as network reach issue or Firewall rules following controller traffic from Gateway.
- To rectify registration failure due to the config error, first update the config settings in the Controller and download the config file. Then shut down Gateway VM, update downloaded config, and then boot Gateway.

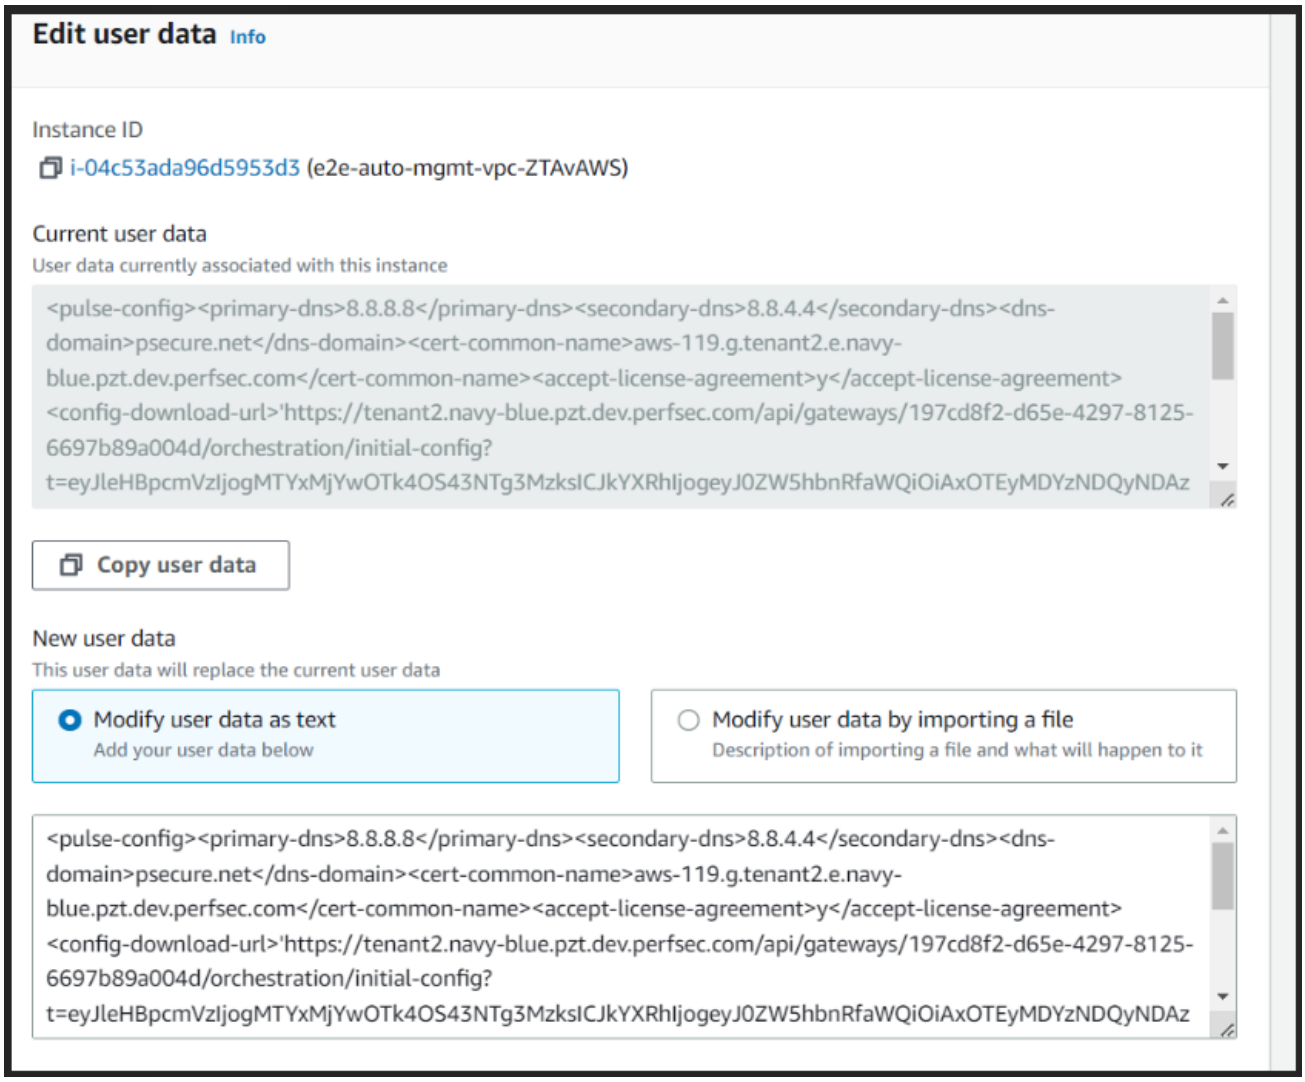

Sample screen: Update Config Value

An option is provided to regenerate and download the gateway init config from the controller admin interface.

Workflow: Creating a Gateway in Microsoft Azure

The process of registering an Azure Gateway with *nZTA* involves two main procedures, to be completed in sequence:

1. Create a Gateway record in the *Controller*.
2. Create a Gateway virtual machine instance in Azure and register it with the *Controller*.

   Azure offers two methods for launching a Gateway virtual machine instance:

   - Through the Azure Marketplace
   - Using the provided template and image files

Before you start, make sure that you have the following information and files for the Gateway:

- An identifying name for the Gateway
- The Gateway geographic location.
- (Optional) The name of the *Gateway Group* to which you want to add this new Gateway record. To learn more about Gateway Groups, see Adding Gateway Groups for High Availability.
- The primary (and optional secondary) DNS server IP address, and search domain.
- The SSH public key that you are using with the Azure Portal or Management Console.

  SSH keys can be generated using sshkeygen on Linux and MacOS, or PuTTyGen on Windows. For further details about generating SSH key pairs, see:

  - For Windows: https://docs.microsoft.com/en-us/azure/virtual-machines/linux/ssh-from-windows
  - For MacOS and Linux: https://docs.microsoft.com/en-us/azure/virtual-machines/linux/mac-create-ssh-keys

- Credentials for the Azure Portal or Management Console.

  These credentials must include sufficient permissions to create a virtual machine.

Additionally, if you are deploying a Gateway instance directly from the template and image files (as opposed to using the Azure Marketplace):

- The Gateway template JSON file:

  *nZTA Gateway*s can be deployed in a new VNET or an existing VNET. Select the JSON template file applicable to your requirements.

  - To deploy in a new VNET:

    https://pulsezta.blob.core.windows.net/gateway/templates/Azure/24-1-371/ivanti-2nic-new-vnet.jsonhttps://pulsezta.blob.core.windows.net/gateway/templates/Azure/24-1-371/ivanti-3nic-new-vnet.json

  - To deploy in an existing VNET:

    https://pulsezta.blob.core.windows.net/gateway/templates/Azure/24-1-371/ivanti-2nic-existing-vnet.jsonhttps://pulsezta.blob.core.windows.net/gateway/templates/Azure/24-1-371/ivanti-3nic-existing-vnet.json

  If you want to use a Management interface, you must download and use the 3 NIC template.

- The Gateway template image VHD file. Choose from:

  - **Americas**: https://pulsezta.blob.core.windows.net/gateway/ISA-V-HYPERV-ZTA-22.7R1-371.1-SERIAL-hyperv.vhd
  - **APJ**: https://pulseztaapj1.blob.core.windows.net/gateway/ISA-V-HYPERV-ZTA-22.7R1-371.1-SERIAL-hyperv.vhd
  - **Europe**: https://pulseztaeurope.blob.core.windows.net/gateway/ISA-V-HYPERV-ZTA-22.7R1-371.1-SERIAL-hyperv.vhd

  Use the link most suitable for your geographic location.

  Registering an Azure Gateway in the *Controller* can also choose to download this file from the **Gateways Overview** page of the *nZTA* Tenant Admin Portal. The opportunity to do this occurs later in this process.

- A public IP address or CNAME for the Gateway. This is the IP address or CNAME at which client devices can externally reach the Gateway instance.

  CNAMEs are not currently supported on *Ivanti Secure Access Client* Linux variants.

Adding an Azure Gateway

To registering a Gateway on your *Controller*, use the **Gateway Details** dialog.

To begin, log into the *Controller* Tenant Admin Portal using the credentials provided in your welcome email. Two outcomes are possible:

- On unconfigured *nZTA* systems, the *Secure Access Setup* Onboarding wizard appears (see Working with the Onboarding Wizard). In this case, click **Add Gateway**.
- On configured *nZTA* systems, the *Network Overview* page appears. In this case:

    1. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways**.

        The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*.

    2. To add a new Gateway, select **Create** from the top-right:

        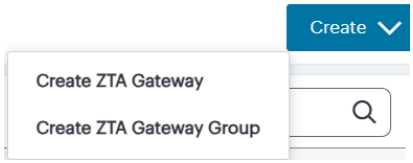

        Add a new Gateway or Gateway Group

    3. In the drop-down menu, click **Create ZTA Gateway**.

In both cases, the **Gateway Details** dialog appears.

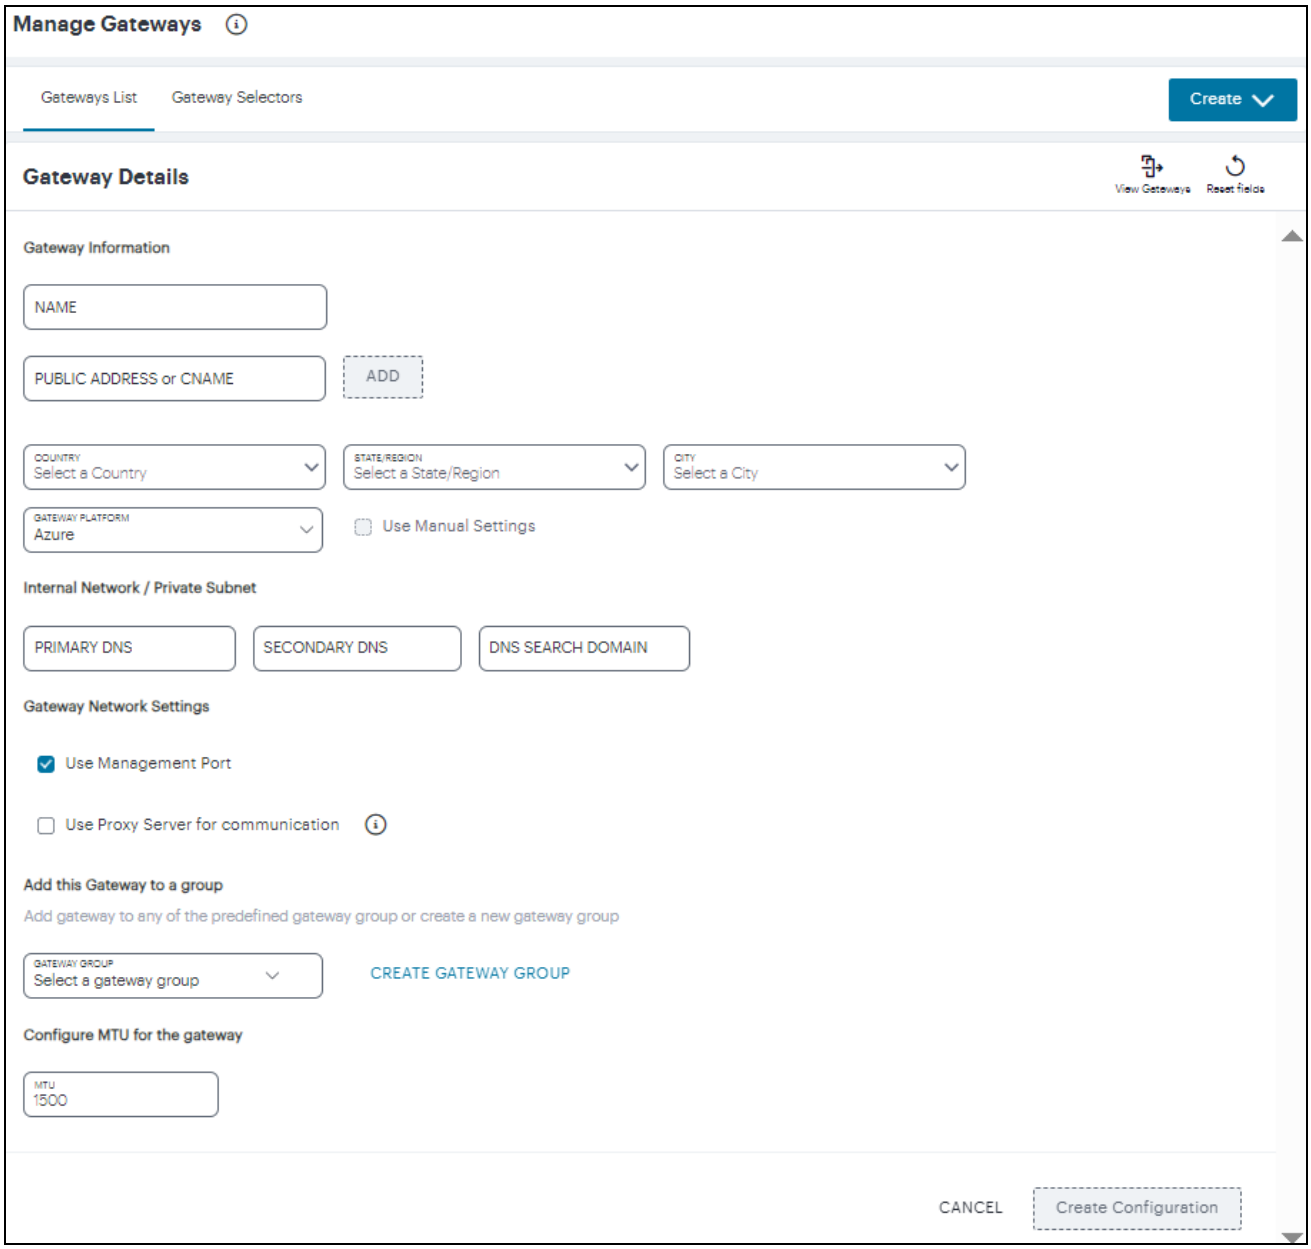

Gateway Details

Enter the following details:

1. Enter a **Name** for the Gateway.
2. Enter one or more **Public Address or CNAME** (Public IP address or CNAME) for the Gateway. Select **Add** to add each entry to the list. To learn more about this setting, see Configuring Networks in your Gateway Datacenter.

   For Azure Marketplace deployments, a public IP address or CNAME is typically allocated at deployment time through the Azure Portal. Therefore, if you do not yet know the expected address/CNAME, enter a dummy value in this field now and update the setting after you have deployed and registered the Gateway instance. For more details on this process, see Creating a Gateway through Azure Marketplace.

3. Select the geographic location details for the Gateway.
4. For **Gateway Platform**, select "Azure".
5. Enter the Primary DNS IP address for the Gateway.
6. (Optional) Enter the Secondary DNS IP address for the Gateway.
7. Enter the DNS Search Domain for the Gateway.

   Make sure the specified DNS service can resolve the IP address of your *Controller*. Issues here can cause registration of the Gateway with the *Controller* to fail.

8. (Optional) Select the **Use Management Port** check box to use management network ports for *nZTA* traffic rather than internal ports.

   When the management port is enabled, Gateway will use management interface to communicate with *Controller* and NTP Server.
   The Gateway will still use the internal port for DNS resolution and NTP server name resolution.
   If the internal DNS cannot resolve the *Controller* domain, the internal interface will require internet access.

9. (Optional) Select the **Use proxy server for communication** check box to enable nZTA to *Controller* communication via proxy server. Proxy is supported on both internal and management interfaces of the Gateway. Once enabled, enter host name and port. Optionally, if your proxy server requires further authentication, enter a username and password to log in to the proxy server.

   Admin can configure proxy for existing Gateway after upgrading it to 22.4R3 version or later.

10. (Optional) Select a Gateway **Group** to which the new Gateway is to be added. To learn more about Gateway Groups, see Adding Gateway Groups for High Availability.

11. **Configure MTU for the gateway**: Configurable MTU size allows admin to modify the default setting of nZTA gateways wherever it is needed. The value allowed is in the range of 576 to 1500 (IPv4).

To add a Gateway definition based on the settings you specified in this dialog, select **Create Configuration**.

After you complete the first part of this workflow, an unregistered Gateway record is created on the *Controller*. This Gateway record can be seen on the **Manage Gateways > Gateways List** page.

Next, create the Gateway virtual machine instance on Azure and allow it to register with the *Controller*. This links the Gateway record on the *Controller* with the actual Gateway virtual machine. Choose to register the Gateway instance through the Azure Marketplace (see Creating a Gateway through Azure Marketplace.) or through the Azure management console (see Creating a Gateway using the Azure Template and Image Files).

Creating a Gateway through Azure Marketplace

This section describes the steps necessary to create the Gateway virtual machine in Microsoft Azure by using the Azure Marketplace. To learn more about the operations included here, refer also to Microsoft's own Azure documentation (https://docs.microsoft.com/azure) for full details.

After you have created an unregistered Gateway record on the **Gateways** page, you must create the Gateway instance in the *Azure portal*. This process facilitates communication between the *Controller* and the Gateway instance.

Before you start, you must obtain the Gateway definition file from the *Controller*. This definition file includes the settings necessary to configure the new Gateway virtual machine with the identity and location of the *Controller*.

To download the Gateway definition file:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways > Gateways List**.

   The *Gateways List* page appears.
3. Locate and select your unregistered Azure Gateway record from the list of available Gateways.
4. Click the **Download** icon to obtain a copy of the Gateway definition file.

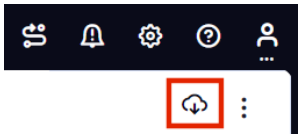

The Download Icon

5. Save the downloaded text file to a location accessible from the Microsoft Azure Portal.

> The Gateway definition file is valid for 24 hours. If this period expires, you must replace the Gateway to generate a new Gateway definition file.

6. (Optional) If you have not yet downloaded the latest version of your Gateway VM file, click the **Download** icon and select **Download gateway VM image**. Save the archive file and unpack to a local workstation. Make sure the resulting file set is accessible from the Microsoft Azure Portal.

To create a Gateway instance from Azure Marketplace:

> *nZTA Gateways* in Azure Marketplace are limited to version 21.3R1 at the present time. To use a Gateway version later than 21.3R1, either launch the Azure Marketplace version and upgrade in-place to the latest version (see Upgrading Gateways) or use the alternate procedure described in Creating a Gateway using the Azure Template and Image Files to launch a Gateway instance using the template and image files.

1. Log into the Microsoft Azure Portal (http://portal.azure.com).
2. Navigate to the Azure Marketplace by clicking **Create a resource**.
3. In the *Search the Marketplace* text box, enter "Ivanti".

   Azure Marketplace presents the results relevant to your search term.
4. Locate *Ivanti Neurons Zero Trust Access Gateway* and click **Create**.

5. In the drop-down list, choose the option that is applicable to your needs:

    - **Ivanti Neurons Zero Trust Access Gateway - BYOL 3 NIC**: Includes 3 network interfaces (internal, external, and management)
    - **Ivanti Neurons Zero Trust Access Gateway - BYOL 2 NIC**: Includes 2 network interfaces (internal and external)

    To first learn more about *Ivanti Neurons Zero Trust Access Gateway*, click the product banner and view the associated information page. You can launch a new Gateway instance from this page.

    The *Create Ivanti Neurons Zero Trust Access Gateway* process appears.

6. On the *Basics* tab, enter the following details:

   - **Subscription**: If you are using the "PZT_Dev" subscription, leave this field as the default value. Otherwise, enter your subscription name.
   - **Resource Group**: Specify the resource group in which the Gateway needs to be deployed, or create a new resource group using the link provided. An Azure Resource Group is a container for a collection of connected assets that you assign to a virtual machine. To learn more, see the Azure documentation (https://docs.microsoft.com/azure).
   - **Region**: Specify the geographic region in which the Gateway instance is deployed.
   - **Ivanti Neurons Zero Trust Access Gateway VM Name**: Enter a suitable name for your Gateway instance. This name must be 1-9 characters long, using only lowercase letters or numbers.
   - **SSH Public Key Source**: Select "Use existing public key".
   - **SSH Public Key**: Copy and paste an RSA public key in a single-line format or the multi-line PEM format.

   SSH keys can be generated using sshkeygen on Linux or MacOS, or PuTTyGen on Windows. For further details about generating SSH key pairs, see:

   - For Windows: https://docs.microsoft.com/en-us/azure/virtual-machines/linux/ssh-from-windows
   - For Linux or MacOS: https://docs.microsoft.com/en-us/azure/virtual-machines/linux/mac-create-ssh-keys

   To continue, click **Next: Network Settings >**.

7. On the *Network Settings* page, enter the following details:

    - **Virtual Network**: A virtual network is a logical isolation of the Azure cloud dedicated to your services. The value you enter here affects the IP address and subnet allocations for all network interfaces shown on this page. Azure pre-populates this field with a new virtual network name, although you can select your own predefined virtual network as necessary.

      To create a new virtual network, perform the following steps:

        1. Click the **Create New** link under the Virtual Network setting.

           The *Create virtual network* dialog appears.

        2. Enter a virtual network name.

        3. Enter an address space in CIDR notation (for example, 192.0.2.0/24).

        4. For each interface subnet, use the automatically-populated name and address values provided, or enter your own details. Each subnet must be contained by the address space entered in the previous setting.

        5. To save your changes, click **OK**.

           Your new virtual network settings are populated into the corresponding interface settings in the main *Network Settings* page.

    - **Internal Subnet**: The subnet identifier for the Internal network, pre-populated by either the selected Virtual Network or your newly-entered virtual network settings.

    - **External Subnet**: The subnet identifier for the External network, pre-populated by either the selected Virtual Network or your newly-entered virtual network settings.

    - (For 3 NIC instances only) **Management Subnet**: The subnet for the Management network, pre-populated by either the selected Virtual Network or your newly-entered virtual network settings.

- **Public IP for Ivanti Neurons Zero Trust Access Gateway external interface LB**: The public IP address identifier at which clients can externally reach the Gateway instance, typically provided by Azure.

  Before you can connect to the new Gateway instance from the *Controller*, you must update the *Controller* with the Public IP address or CNAME assigned to the external interface of the Gateway load balancer. This process is described later.

- **DNS prefix for external interface LB**: The unique DNS name for the public IP address specified for the external interface load balancer.
- **Public IP for NAT Gateway**: The public IP address identifier of a NAT Gateway for the virtual machine to communicate with the *Controller* and other public resources.
- **DNS prefix for NAT Gateway public IP**: The unique DNS name for the public IP address specified for the internal interface NAT Gateway.
- **Deploy Ivanti Neurons Zero Trust Access Gateway with Load Balancer**: To deploy this Gateway with a load balancer, select "Yes" from the drop-down list. The front-end IP address of the load balancer is then used as the public IP address for your Gateway.

  If you select "No" to not deploy a load balancer, you must create and associate a public IP address to the external interface of your instance after deployment is complete.

  In all cases, on completion of this process, you must update the *Controller* Gateway definition with the correct public IP address for your Azure Gateway instance.

To continue, click **Next: Instance Configuration >**.

8. On the *Instance Configuration* page, enter the following details:

    - **Ivanti Neurons Zero Trust Access Gateway VM Size**: This is the specification of the virtual machine. Choose from:

        - For 2nic instances, select "1 x Standard DS2 v2"
        - For 3nic instances, select "1 x Standard DS4 v2"

    - **Diagnostic storage account**: The storage account for the virtual machine diagnostics. The default value is a new account based on your VM name.
    - **Ivanti Neurons Zero Trust Access Gateway Version**: Specify the version applicable to the current *nZTA* version, or the version you require. *Ivanti* recommends you select the latest available version.
    - **Ivanti Neurons Zero Trust Access Gateway Config Data**: Paste in the raw text of your Gateway definition file. To obtain the Gateway definition file, see the process described earlier in this section.

    To continue, click **Next: Review + create >**.

9. On the *Review + create* page, verify the proposed configuration is validated successfully, and then click **Create** to create your new Gateway instance.

    After a short wait, your instance is created and deployed.

10. Access the virtual machine settings for your new Gateway instance, and click **Networking** from the *Settings* menu.

    The *Networking* dialog appears, showing your attached network interfaces (internal, external, and (optionally) management).

11. Click the tab that corresponds to the *external* network interface.

    The settings for the external network interface appear.

12. Locate the **NIC Public IP** field and make a note of the IP address shown there. This is the public IP address you use to reconfigure the *Controller* record for this Gateway.

    If no public IP address is shown, determine if a load balancer was deployed together with your Gateway instance by selecting the **Load balancing** tab.

    - If a load balancer was deployed, make a note of the **Frontend IP address** displayed in this tab and use this as the Gateway public IP address on the *Controller*.
    - If a load balancer was not deployed, create a public IP address and associate it with the *external* interface. Then, use this IP address as the Gateway public IP address on the *Controller*.

    > To learn about configuring IP addresses in the Azure portal, see the Microsoft Azure documentation.

13. Return to the *nZTA* Tenant Admin Portal, and click **Secure Access > Gateways > Gateways List**.

    The *All Gateways* page appears.

14. Select your new Azure Gateway from the list.

    The *Gateways Overview* page appears.

15. Make sure the new Azure Gateway instance is shown in the list of configured Gateways and is connected (Status is *Online* and State is *Registered*).

16. Select **Secure Access > Manage Gateways > *Gateway* > Configuration**, and locate *Gateway Network Settings*. Enter the Public IP address you noted from the Azure virtual machine settings. Make sure you remove any previously-entered dummy values.

17. To save your changes, click **Save Changes**.

This completes the Azure Gateway registration process. Your enrolled client devices should now be able to connect to the Gateway.

You can view all Gateways that are deployed in your *nZTA* service (see Viewing and Monitoring Gateways in the Controller) or you can move to the next stage of *nZTA* configuration, which is Working with User Authentication.

Creating a Gateway using the Azure Template and Image Files

This section describes the steps necessary to create the Gateway virtual machine in Microsoft Azure by using the provided template and image files. To learn more about the operations included here, refer also to Microsoft's own Azure documentation for full details.

After you have created an unregistered Gateway record on the **Gateways** page, you must create the Gateway instance in the *Azure Management Console*. This process facilitates communication between the *Controller* and the Gateway instance.

Before you start, you must complete the following prerequisites:

- Obtain the Gateway definition file from the *Controller*. This definition file includes the settings necessary to configure the new Gateway virtual machine with the identity and location of the *Controller*.

To download the Gateway definition file:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways > Gateways List**.

   The *Gateways List* page appears.
3. Locate and select your unregistered Azure Gateway record from the list of available Gateways.
4. Click the **Download** icon to obtain a copy of the Gateway definition file.

   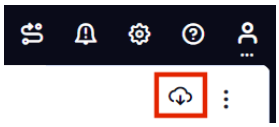

   The Download Icon
5. Save the downloaded text file to a location accessible from the Microsoft Azure Console.

The Gateway definition file is valid for 24 hours. If this period expires, you must replace the Gateway to generate a new Gateway definition file.

1. (Optional) If you have not yet downloaded the latest version of your Gateway VM file, click the **Download** icon and select **Download gateway VM image**. Save the archive file and unpack to a local workstation. Make sure the resulting file set is accessible from the Microsoft Azure Portal.

- Create a new Azure *Resource Group* in your desired location and subscription account.
- Create a new *storage account*, and create a new container in that account.

- Download the *nZTA* Azure VHD image file for your region and copy it to the storage account you created in the previous step.

  Choose to download from the following regions:

  - **Americas**: https://pulsezta.blob.core.windows.net/gateway/ISA-V-HYPERV-ZTA-22.7R1-371.1-SERIAL-hyperv.vhd
  - **APJ**: https://pulseztaapj1.blob.core.windows.net/gateway/ISA-V-HYPERV-ZTA-22.7R1-371.1-SERIAL-hyperv.vhd
  - **Europe**: https://pulseztaeurope.blob.core.windows.net/gateway/ISA-V-HYPERV-ZTA-22.7R1-371.1-SERIAL-hyperv.vhd

  For this process, you can use *azcopy*:

  1. From the storage account, create a Shared Access Signature (SAS) token:

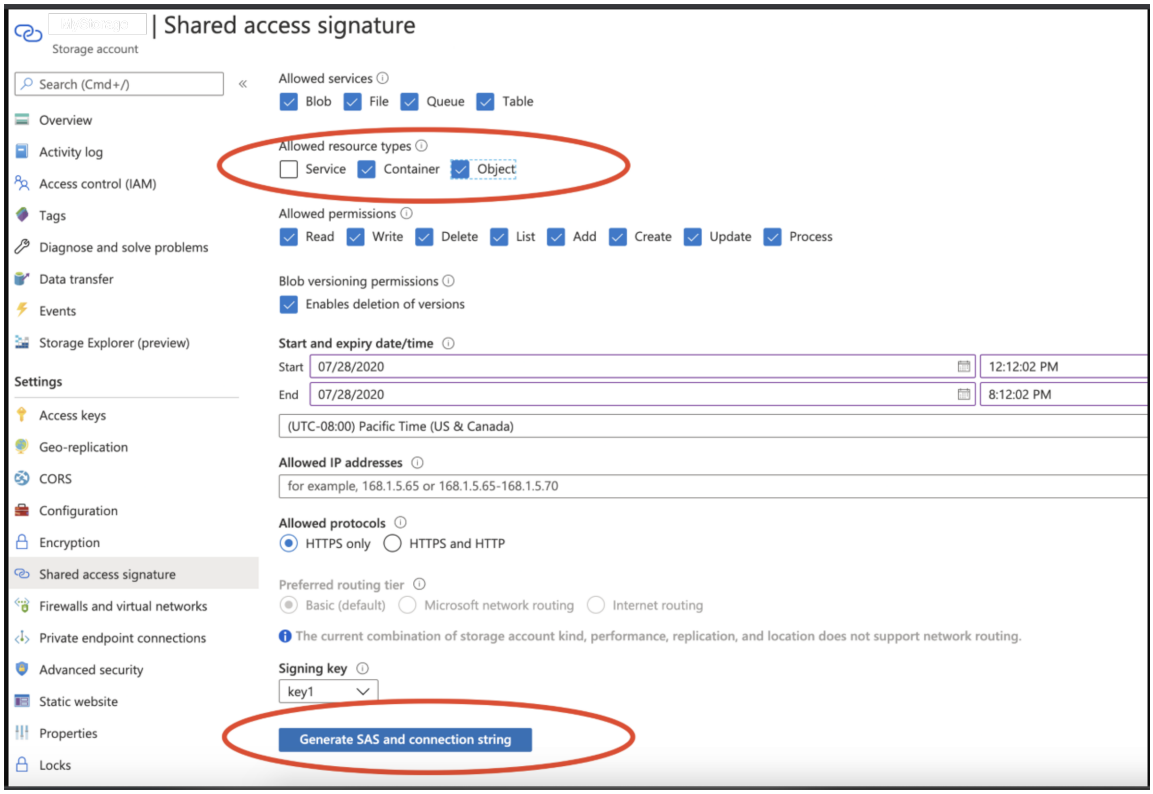

Creating a SAS token from a storage account in Azure

2. Open the Azure Cloud Shell and start a bash shell:

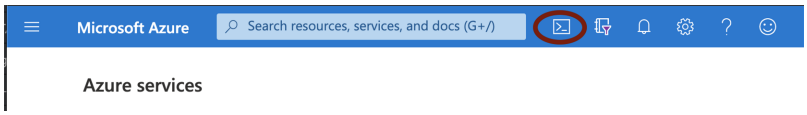

Starting the Azure Cloud Shell

3. In the Azure Cloud shell, use *azcopy* to copy the Gateway VHD image file into your storage account. For example, use the following syntax:

```
azcopy copy '<URL to VHD file>'
'https://<MyStorageAccount>.blob.core.windows.net/<Container_Name>/<VHD
filename><SAS-Token>''
```

Replace the angled-bracket elements with the details gathered in previous steps. For example:

```
azcopy copy 'https://pulseztaeurope.blob.core.windows.net/gateway/ISA-V-HYPERV-
ZTA-22.7R1-371.1-SERIAL-hyperv.vhd'
'https://MyStorage.blob.core.windows.net/gateway/ISA-V-HYPERV-ZTA-22.7R1-371.1-
SERIAL-hyperv.vhd?sv=2024-03-23&ss=bfqt&srt=co&sp=rwdlacupx&se=2024-03-
23T02:57:39Z&st=2023-09-
28T18:57:39Z&spr=https&sig=mJU7WNd9oNY7wcXNOqEOhbYshD9Sxv56rqEl%2FmEuC
g4%3D'
```

To create a Gateway instance using the template and image files specified in the release notes for this product version:

For reference, the recommended minimum requirements for a Gateway virtual machine instance in Azure are:

- Standard_D2s_v3 (2 vCPU, 8 GB Memory), or
- Standard_F4s (4 vCPU, 8 GB Memory)

1. Access the Azure Management Console and log in using your credentials.
2. Access the **Home > Templates** page to view available templates.
3. Click "+ Add" to add a new template.
4. In the new template "General" section, enter a template name and description.

5. In the new template "ARM Template" section, remove the default data and replace with the raw text contents of the *nZTA* Azure Gateway template JSON file.

   > Use either the *new VNET* template JSON file or the *existing VNET* template JSON file as per your requirements.

6. Save the new template.
7. On the **Home > Templates** page, locate the new Azure Gateway template.
8. On the context menu for the template, click **Deploy**.

   The **Custom Deployment** page appears.

9. On the **Custom Deployment** page, enter any required details for the Gateway deployment.

   - **Resource Group**: Specify the resource group name in which the Gateway needs to be deployed, or create a new group.
   - **Location**: Specify the region in which the Gateway instance is deployed.
   - **ZTAStorage Account Name**: Specify the storage account you created earlier where the Gateway image is held.
   - **ZTAStorage Account Resource Group**: Specify the resource group you created earlier.
   - **ZTAImage Location URI**: Enter the full URI for the Gateway template image VHD file you copied to your storage account earlier.
   - **ZTAVM Name**: Enter a suitable name for your Gateway instance. *Ivanti* recommends matching the Gateway name used during the process of creating the Gateway record on the *Controller*.
   - **ZTAConfig**: Paste in the raw text of your Gateway definition file. To obtain the Gateway definition file, see the process described earlier in this section.
   - **SSH Public Key**: Specify your SSH key pair name.

10. If required, update the labels for the instance:
    - Update the **Dns Label Prefix Ext** if required. For example, "azuregwext".
    - Update the **Dns Label Prefix Mgmt** if required. For example, "azuregwmgmt".
    - Update the **Existing Vnet Name** if required.
    - Update the **Existing Internal Subnet** if required. For example: "InternalNW".
    - Update the **Existing External Subnet** if required. For example: "ExternalNW".
    - Update the **Existing Management Subnet** if required. For example: "ManagementNW".

11. For **Load Balancer Configuration**, If you plan to deploy multiple Gateways inside a Gateway Group, select "Yes" to deploy a new internet-facing Public Standard Load Balancer instance alongside the Gateway. Select "No" to launch only this Gateway instance.

    This option is applicable only for new VNET templates.

    If you elect to launch a load balancer, the following pre-configuration is applied:

    - A Standard SKU Public IP address is assigned to the Load Balancer.
    - A Backend Pool is created and the deployed Gateway is associated with the pool through it's external network interface.
    - A health probe is configured on TCP port 443.
    - Load balancing rules are configured.

    After you have deployed the Gateway and Load Balancer, you must return to the Tenant Admin Portal on the *Controller* and update the Gateway Group **Load Balancer IP ADDRESS** setting to be the Load Balancer's public IP address.

    If you want to configure the Load Balancer to balance across further Gateway instances from the Gateway Group, you must deploy each subsequent Gateway into the same Resource Group through the use of existing VNET templates and then update the Load Balancer's Backend Pool.

    With new VNET templates, a NAT gateway is deployed for routing outbound Internet traffic from the Gateway's internal network interface in order for the Gateway to be able to reach the *Controller*.

    Public IP addresses are not automatically assigned to any of your Gateway's network interfaces. If you are deploying a Gateway into an existing VNET, in order for the Gateway to be able to reach the *Controller* from it's internal network interface, make sure you allow outbound Internet traffic from the Private Subnet for the deployed Gateway.

    To learn more about high availability and Gateway Groups, see Adding Gateway Groups for High Availability and Using Gateway Groups for High Availability.

12. Agree to the terms and conditions.

13. Click **Purchase** to start the creation of the Gateway.

    A window displays the status of the process, starting with **Deployment in Progress**.

    (Optional) Click the **Deployment in Process** hyperlink to view a status page for the process.

14. Wait until the process completes.

15. Ensure that your Azure Security Groups support the IP addresses allocated to the Gateway instance. Please refer to Azure's own documentation for full details.

16. *(This step is required only if you have not deployed your Gateway with a Load Balancer or NAT at the front-end)* Public IP addresses are not automatically assigned to any of the Gateway's network interfaces. Therefore, before your client devices can connect to the new Gateway instance from the *Controller*, you must associate a new Public IP address with the external interface of the Gateway. Then, update the *Controller*'s Gateway Public IP address setting to match this address (in the **Secure Access > Gateways > Gateways List** page, select your new Gateway, then select **Secure Access > Gateways > Configuration** and locate *Gateway Network Settings*).

17. In the **Secure Access > Gateways > Gateways List** page, make sure the new Gateway has a confirmed status of *Connected*.

    This completes the Azure Gateway registration process. Your enrolled client devices should now be able to connect to the Gateway.

18. (Optional) After you have registered a Gateway, you can configure it (or the Gateway Group to which it belongs) as the default Gateway if required. A default Gateway handles all requests from applications that are not referenced by any secure access policy. See Configuring a Default Gateway for Application Discovery.

You can view all Gateways that are deployed in your *nZTA* service (see Viewing and Monitoring Gateways in the Controller.) or you can move to the next stage of *nZTA* configuration, which is Working with User Authentication.

Workflow: Creating a Gateway in KVM/OpenStack

This workflow leads you through the processes for setting up a KVM Gateway in OpenStack. These processes must be performed in sequence:

- Preparing to create a KVM gateway, see Preparing to Create a KVM Gateway.
- Creating the gateway record in the *Controller*, see "Adding a KVM Gateway" on page 288.

- Preparing Metadata for OpenStack, see Preparing Metadata for OpenStack.
- Creating the KVM Gateway virtual machine instance in OpenStack, see Creating the KVM Gateway Virtual Machine Instance in OpenStack.

After these steps have been completed successfully, the *Controller* and Gateway establish communication with each other.

Preparing to Create a KVM Gateway

Before you start, make sure you have the following information and files:

- An identifying name for the Gateway
- The public IP address for the Gateway. This is the IP address at which clients can externally reach the Gateway instance.
- The Gateway geographic location
- (Optional) The name of the *Gateway Group* to which you want to add this new Gateway record. To learn more about Gateway Groups, see the *Tenant Admin Guide*.

Additionally, to manually specify KVM Gateway network interface settings:

- The primary (and optional secondary) DNS server IP address, and search domain.
- The required internal/private subnetworks must already be defined on OpenStack. Please refer to the OpenStack documentation for details.
- The required external subnetworks must already be defined on OpenStack. Please refer to the OpenStack documentation for details.
- *(Optional)* Any required management subnetwork must already be defined on OpenStack. Please refer to the OpenStack documentation for details.

- The Gateway KVM template: https://pulsezta.blob.core.windows.net/gateway/ISA-V-KVM-ZTA-22.7R1-371.1.zip

  Download a copy of the KVM template ZIP file. Then fully unpack the ZIP file, including any compressed *.gz* files inside, to a local workstation. Make sure that the resulting file set is accessible from the OpenStack Console.

  You can also choose to download this file from the **Gateways Overview** page of the *nZTA* Tenant Admin Portal. The opportunity to do this occurs later in this process.

- Credentials for the OpenStack Console.

  These credentials must include sufficient permissions to create a virtual machine from a template image.

After you have all required information, you can set up a *nZTA* KVM gateway, see "Adding a KVM Gateway" below.

Adding a KVM Gateway

To registering a Gateway on your *Controller*, use the **Gateway Details** dialog.

To begin, log into the *Controller* Tenant Admin Portal using the credentials provided in your welcome email. Two outcomes are possible:

- On unconfigured *nZTA* systems, the *Secure Access Setup* Onboarding wizard appears (see Working with the Onboarding Wizard). In this case, click **Add Gateway**.

- On configured *nZTA* systems, the *Network Overview* page appears. In this case:

    1. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways**.

       The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*.

    2. To add a new Gateway, select **Create** from the top-right:

       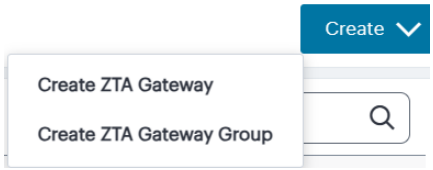

       Add a new Gateway or Gateway Group

    3. In the drop-down menu, click **Create ZTA Gateway**.

The **Gateway Details** dialog appears.

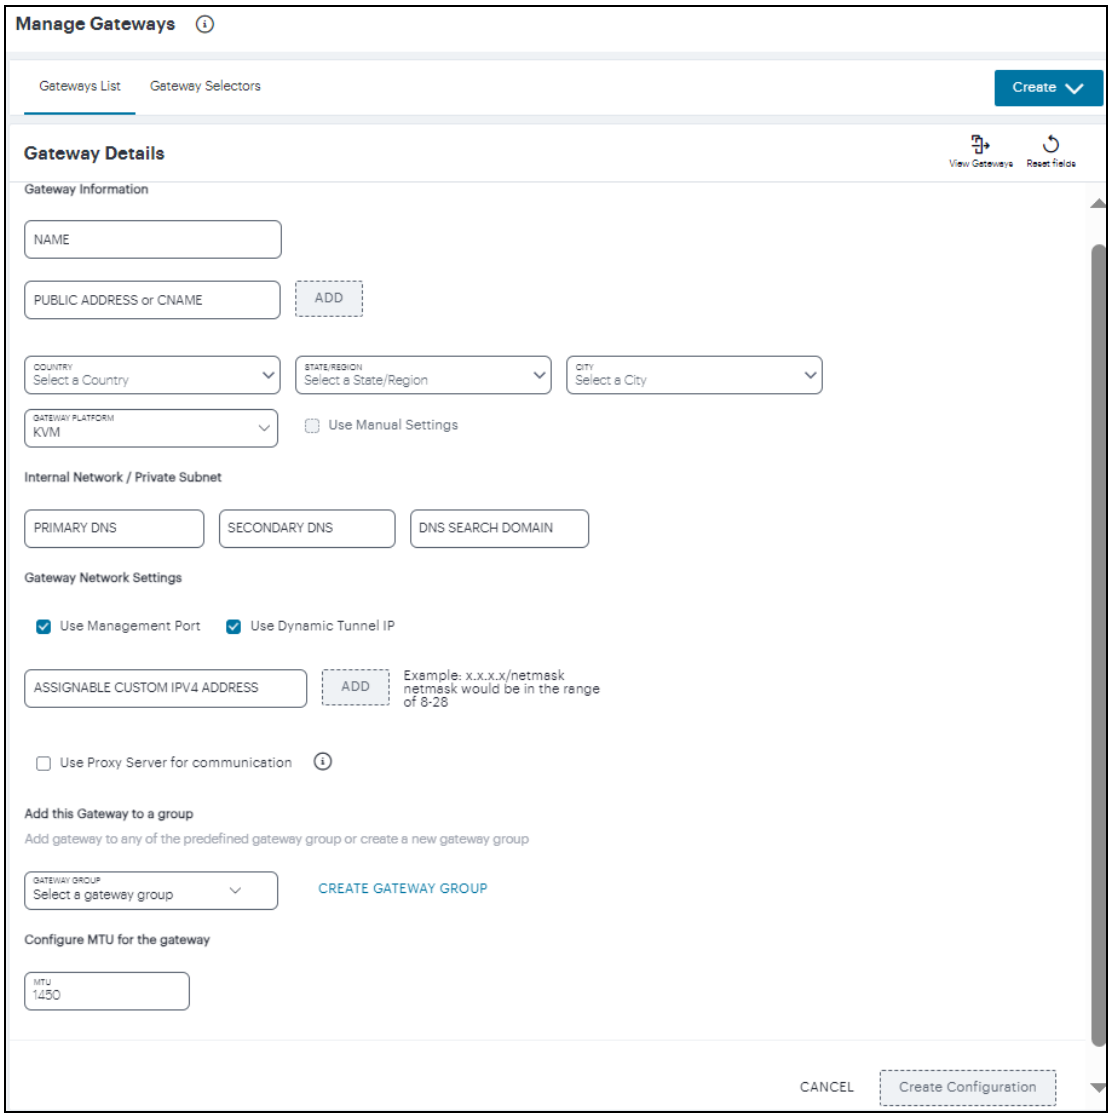

Gateway Details

Enter the following details:

1. Enter a **Name** for the Gateway.
2. Enter one or more **Public Address or CNAME** (Public IP address or CNAME) for the Gateway. Select **Add** to add each entry to the list. To learn more about this setting, see Configuring Networks in your Gateway Datacenter.
3. Select the geographic location details for the Gateway.
4. For **Gateway Platform**, select "KVM".

5. Enter the Primary DNS IP address for the Gateway.
6. (Optional) Enter the Secondary DNS IP address for the Gateway.
7. Enter the DNS Search Domain for the Gateway.

> Make sure the specified DNS service can resolve the IP address of your *Controller*. Issues here can cause registration of the Gateway with the *Controller* to fail.

8. (Optional) Select the **Use Management Port** check box to use management network ports for *nZTA* traffic rather than internal ports.

> When the management port is enabled, Gateway will use management interface to communicate with *Controller* and NTP Server.
> The Gateway will still use the internal port for DNS resolution and NTP server name resolution.
> If the internal DNS cannot resolve the *Controller* domain, the internal interface will require internet access.

9. (Optional) Select the **Use Dynamic Tunnel IP** check box to configure a pool of IP addresses that are dynamically mapped to client sessions with this Gateway, such that user traffic from the Gateway to an application can be identified as originating from a specific client. To learn more, see Using Dynamic IP Addressing to Profile Client Traffic.

   The *Custom IP Pool* dialog appears:

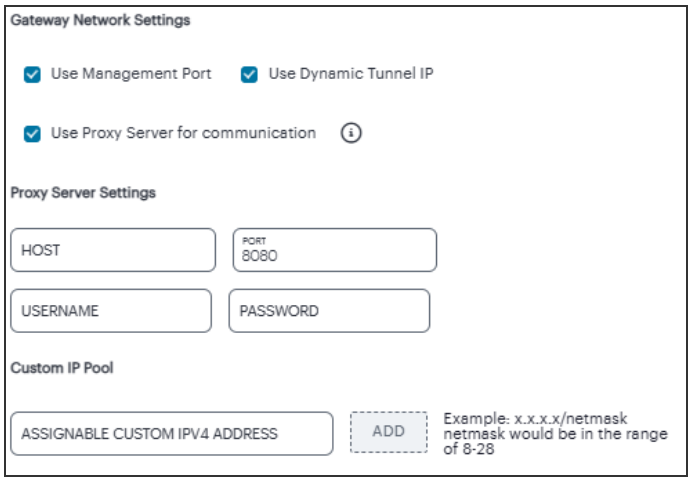

   Gateway Details - Custom IP Pool settings

   Dynamic Tunnel IP addresses are not supported in Gateway Groups.

   Use the **Assignable Custom IPv4 Address** field to enter an IP address and subnet (in the range 8-28) in CIDR notation, then click **Add**. Repeat this step for each address/subnet you want to use.

10. (Optional) Select the **Use proxy server for communication** check box to enable nZTA to *Controller* communication via proxy server. Proxy is supported on both internal and management interfaces of the Gateway. Once enabled, enter host name and port. Optionally, if your proxy server requires further authentication, enter a username and password to log in to the proxy server.

    Admin can configure proxy for existing Gateway after upgrading it to 22.4R3 version or later.

11. (Optional) Select a Gateway **Group** to which the new Gateway is to be added. To learn more about Gateway Groups, see Adding Gateway Groups for High Availability.

12. **Configure MTU for the gateway**: Configurable MTU size allows admin to modify the default setting of nZTA gateways wherever it is needed. The value allowed is in the range of 576 to 1450 (IPv4).

13. To add a Gateway definition based on the settings you specified in this dialog, select **Create Configuration**.

After you complete the first part of this workflow, an unregistered Gateway record is created on the *Controller*. This Gateway record can be seen on the **Gateways > Gateways List** page.

You can now prepare your metadata, see Preparing Metadata for OpenStack.

Preparing Metadata for OpenStack

The preparation of metadata for use on OpenStack currently requires some manual steps:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.

2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways > Gateways List**.

   The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*.

3. locate and select your KVM Gateway.

   The *Gateways Overview* page appears.

4. Click the **Download** icon to obtain a copy of the Gateway definition file.

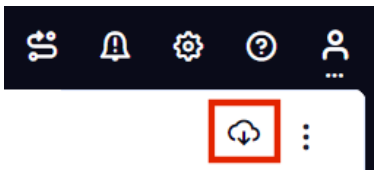

The Download Icon

5. Specify a save location for your Gateway definition file.

> The Gateway definition file is valid for 24 hours. If this period expires, you must replace the Gateway to generate a new Gateway definition file.

6. (Optional) If you have not yet downloaded the latest version of your Gateway VM file, click the **Download** icon and select **Download gateway VM image**. Save the archive file and unpack to a local workstation. Make sure the resulting file set is accessible from the OpenStack Management Portal.
7. View the gateway definition file in a text editor.

8. Start a separate text editor file, and paste the following template text block into it:

```
<pulse-config>
  <config-download-url>
     '<insert vaConfigURL value here>'
  </config-download-url>
  <appliance-id>
     <insert vaApplianceID value here>
  <secondary-dns>
     <insert vaSecondaryDNS here>
  </secondary-dns>
  <primary-dns>
     <insert vaPrimaryDNS here>
  </primary-dns>
  <dns-domain>
     <insert vaDnsSearchDomain here>
  </dns-domain>
  </appliance-id>
  <cert-common-name>
     <insert vaCommonName value here>
  </cert-common-name>
  <accept-license-agreement>
     y
  </accept-license-agreement>
  <controller-enrolled-hostname>
     <insert va|CTR_name|EnrolledHostname value here>
  </controller-enrolled-hostname>
  <dns-search-domain>
     <insert vaDnsSearchDomain value here>
  </dns-search-domain>
  <controller-hostname>
     <insert va|CTR_name|Hostname value here>
  </controller-hostname>
</pulse-config>
```

9. For each parameter block in the template text block file:

    - Locate the required metadata property for the line.

      For example, in the following block:

      ```
      <appliance-id>
         <insert vaApplianceID value here>
      </appliance-id>
      ```

      You require the **vaApplianceID** value from the gateway definitions file.

    - Locate the required value in the gateway definitions file.

      For example, the **vaApplianceID** value is *99ce3aa3c9494cbabb51c085c9c3f6ad*.

    - Copy and paste this value from the gateway definitions file into the template text file.

      For example, the <appliance-id> block will now read as follows:

      ```
      <appliance-id>
         99ce3aa3c9494cbabb51c085c9c3f6ad
      </appliance-id>
      ```

    > You do not need to change the <accept-license-agreement> block, and can retain its *y* setting.

10. After you have added all required text to the template text file, save that file for use in the next section.

You can now create a KVM gateway VM in Openstack, see Creating the KVM Gateway Virtual Machine Instance in OpenStack.

Creating the KVM Gateway Virtual Machine Instance in OpenStack

To create a KVM VM instance in OpenStack:

1. Access the *OpenStack Management Portal*, either from a client or a web browser, and log in using your OpenStack credentials.

   In the OpenStack console, the **Overview** page appears.

2. In the left menu, click **Compute > Images**.

   The **Images** page appears. This shows a list of images.

3. Above the list of images, click **Create Image**.

   The **Create Image** wizard appears. In this wizard, you upload a KVM gateway image for use.

4. Under **Image Details**:

   - Enter an **Image Name**. Typically, this incorporates a version number. For example, *ZTA_GWY_100*.
   - Enter an **Image Description**. For example: *ZTA KVM Image*.
   - Under **Image Source**, click **Browse** and select the unpacked KVM disc image file. Then, click **Format** and select *QCOW2 - QEMU Emulator*.
   - Under **Image Requirements**, set **Minimum Disk (GB)** to *40* and **Minimum RAM (KB)** to *2048*.
   - Set **Visibility Setting** as required. *Public* will enable the image to be used in other projects. *Private* will not.
   - Set **Image Sharing** as required.
   - Use the default settings for all other properties.

5. Click **Next**.

   The **Metadata** page of the wizard appears. No action is required on this page, all properties can use their default settings.

6. Click **Create Image**.

   The wizard closes, and the new KVM gateway image is added to the **Images** page.

7. Wait until the image has been uploaded and processed and shows as *Active*.

   The upload image process typically takes 15-20 minutes.

8. After the image has uploaded and is *Active*, click its **Launch** button.

   The first page of the **Launch Instance** wizard appears. In this wizard, you create a KVM gateway instance.

9. Under **Details**:

   - Enter an **Instance Name**. This will be the displayed name of the gateway in *nZTA*.
   - Enter a **Description** for the KVM gateway. For example *ZTA KVM Gateway*.
   - Use the default settings for all other properties.

10. Click **Next**.

    The **Source** page of the wizard appears. This page lists the selected disk image and selected/default settings for the instance. No action is required on this page, all properties can use their displayed settings.

11. Click **Next**.

    The **Flavor** page of the wizard appears. This page lists the available types of gateway you can create.

12. Locate the *ISA4000-V* entry and click its "up arrow" button to select it.

13. Click **Next**.

    The **Networks** page of the wizard appears. This page lists the available networks (and associated subnetworks) for the gateway. It enables you to select the required subnetworks for your gateway.

14. In the available list, locate the required subnetworks.

    For example, you may require a subnetwork for internal ports and a subnetwork for external ports, but not a subnetwork for management interfaces.

    

    If the required subnetworks do not yet exist, you must define them. Please refer to the OpenStack documentation for details of this process.

15. Click the "up arrow" button for each subnetwork to select it.

    For each selected subnetwork, a fixed IP address is added automatically to the gateway. These appear later in this process, so that they can be assigned to floating IP addresses.

16. Click **Next**.

    The **Network Ports** page of the wizard appears. No action is required on this page, all properties can use their displayed settings.

17. Click **Next**.

    The **Security Groups** page of the wizard appears. No action is required on this page, all properties can use their displayed settings.

    If there is no default security group defined, you must define one. Please refer to the OpenStack documentation for details of this process.

18. Click **Next**.

    The **Key Pair** page of the wizard appears. No action is required on this page, all properties can use their displayed settings.

19. Click **Next**.

    The **Configuration** page of the wizard appears. This page enables you to configure the gateway instance using metadata you prepared earlier, see Preparing Metadata for OpenStack.

20. Open your template text file and copy the entire text block that starts with `<pulse-config>` and ends with `</pulse-config>`.

21. Paste the text block into the **Customization Script** block.

    You cannot directly paste metadata for your gateway from *nZTA*. You must prepare a suitable text block from the metadata, see Preparing Metadata for OpenStack.

22. Enable the **Configuration Drive** check box.

23. Click **Launch Instance**.

    The wizard closes, and the new KVM gateway instance is added.

24. Access the **Instances** page.

    The new KVM gateway instance is listed on this page.

25. Wait until the **Power State** of the gateway instance is *Running*.

    > This process may take several minutes.

26. After the instance state changes to *Running*, make a note of the subnetworks and their automatically-assigned fixed IP addresses in the **IP Address** column for the instance. For example:

    

    Unassociated KVM Ports

    In this example, floating IP addresses are listed after the fixed IP addresses, so all are unassociated:

    - The fixed IP address on the *int-port-2* subnetwork is *4.4.10.83*.
    - The fixed IP address on the *ext-port-2* subnetwork is *5.5.10.64*.

27. Access the **Network > Floating IPs** page.

    The **Floating** IPs page shows the floating IP addresses associated with your account. Both associated and unassociated floating IP addresses are listed.

    > Associated floating IPs have a **Mapped Fixed IP Address** listed.

28. Identify an unassociated floating IP address that you want to associate with a fixed IP address.

29. Click the **Associate** button for the fixed IP address.

    The **Manage Floating IP Associations** dialog appears.

30. Select a fixed **port to be associated** for the selected floating IP address.

31. Click **Associate** to conform the association.

32. Repeat the association process until each of the fixed IP addresses for your gateway instance is associated with a floating IP address.
33. Wait until the status of these floating IP addresses all show as *Active*.
34. Return to the **Compute** > **Instances** page.

    This page now shows a fixed IP address associated with floating IP address for each port. For example:

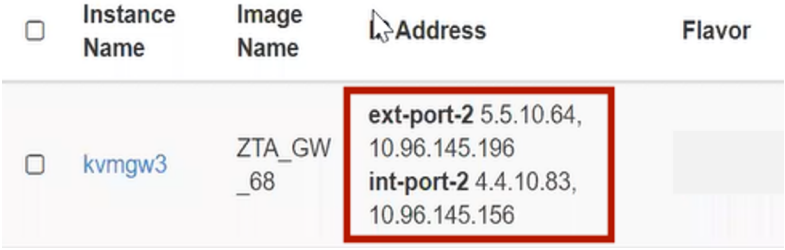

    Associated KVM Ports

35. Click the **Console** tab.

    A console monitor view shows the ongoing boot-up process for the instance.

36. Wait until the instance shows a screen similar to the following:

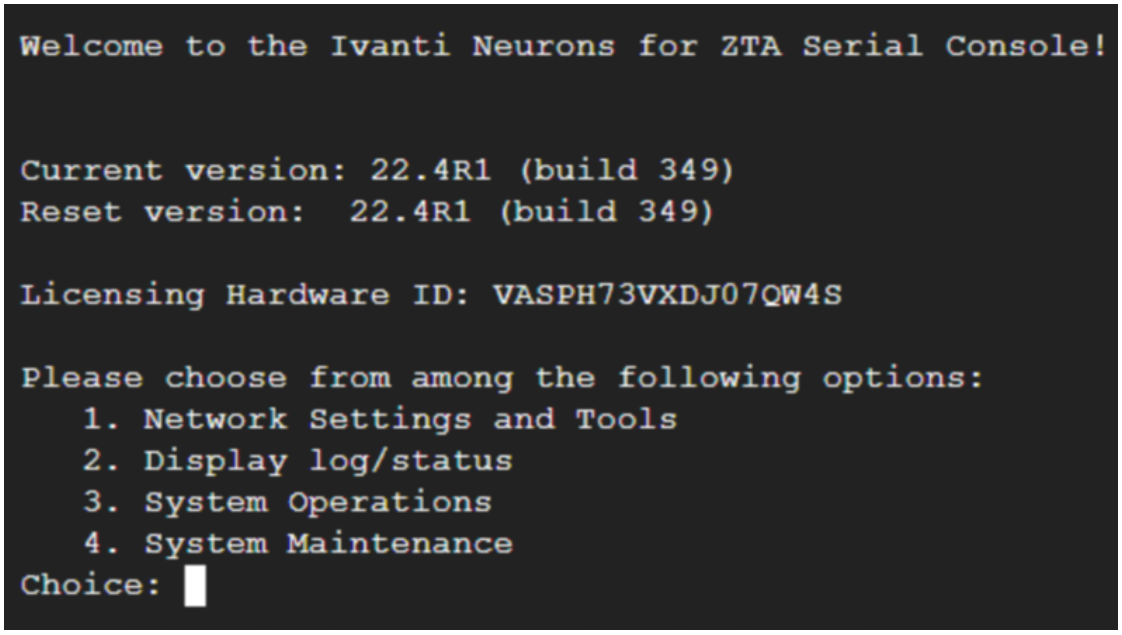

    KVM Instance Console Monitor

37. Return to the **Gateways List** page on the *Controller*.

38. Locate the new Gateway record in the list and confirm that its status has updated to *Connected*. For example:

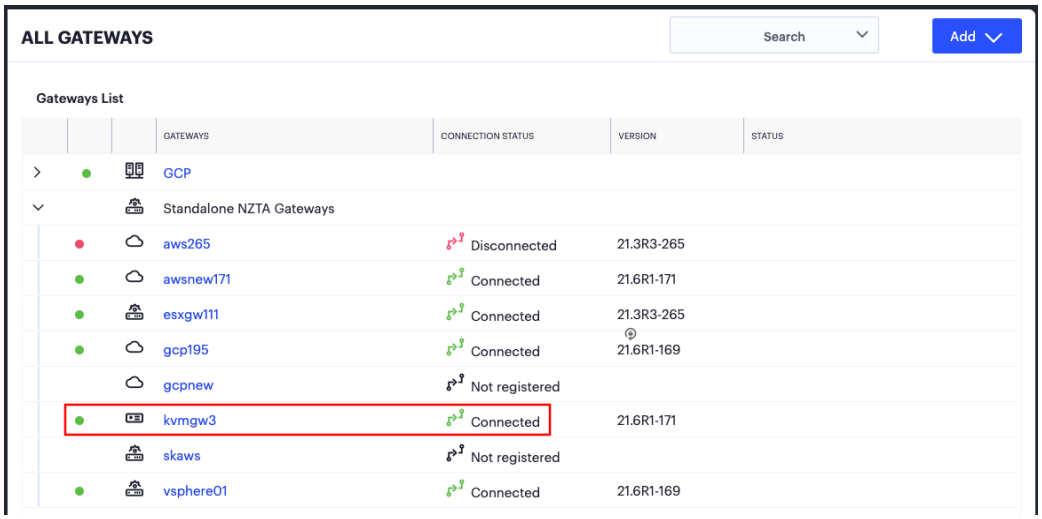

KVM Gateway Connected

39. (Optional) After you have registered a Gateway, you can configure it (or the Gateway Group to which it belongs) as the default Gateway if required. A default Gateway handles all requests from applications that are not referenced by any secure access policy. See Configuring a Default Gateway for Application Discovery.

Workflow: Creating a Gateway in Google Cloud Platform

This workflow leads you through the processes for setting up a Gateway on the Google Cloud Platform (GCP). These processes must be performed in sequence:

- Preparing to create a GCP gateway, see Preparing to Create a GCP Gateway.
- Creating the gateway record in the *Controller*, see "Adding a GCP Gateway" on page 307.
- Downloading Metadata for Google Cloud Platform, see Downloading Metadata for Google Cloud Platform.
- Uploading the GCP Image onto the Google Cloud Platform, see Uploading the GCP Virtual Machine Image onto the Google Cloud Platform.

- Creating a VM Instance of the GCP image. Either:
  - Creating a VM Instance of the Uploaded GCP Image Manually, see Creating a VM Instance of the Uploaded GCP Image Manually.
  - Creating a VM Instance of the Uploaded GCP Image Using a Script/Template, see Creating a VM Instance of the Uploaded GCP Image Using a Script/Template.
- Completing the Configuration of the *Controller*, see Completing the Configuration of the Controller.

After these steps have been completed successfully, the *Controller* and Gateway establish communication with each other.

Preparing to Create a GCP Gateway

Before you start, make sure you have the following information and files:

- An identifying name for the Gateway.
- The public IP address for the Gateway. This is the IP address at which clients can externally reach the Gateway instance, such as an LB/NAT or Datacenter network forward rules.

  If you want Google Cloud platform to allocated a public IP address automatically, you can use a dummy IP address (for example, *1.1.1.1*) when you create the Gateway on *nZTA*. You must then update the *Controller* with the allocated public IP address afterwards.

- The Gateway geographic location.
- (Optional) The name of the *Gateway Group* to which you want to add this new Gateway record. To learn more about Gateway Groups, see "Adding Gateway Groups for High Availability" on page 238

  Gateway Group may have a defined public IP address, which you can specify during the creation of the Gateway.

- The *nZTA Gateway* GCP virtual machine image: https://pulsezta.blob.core.windows.net/gateway/ISA-V-GCP-ZTA-22.7R1-371.1.tar.gz

  Download a copy of the GCP Gateway image as a compressed TAR archive file, then decompress the archive to a local workstation. Make sure that the resulting file set is accessible from the Google Cloud Platform Console.

  > You can also choose to download the Gateway image through the **Gateways Overview** page of the *Controller* after you have defined the Gateway record. The opportunity to do this occurs later in this process.

- *(Optional)* GCP Gateway YAML templates, suitable for automating the creation of your GCP VM instances. Choose from:

  - To deploy in an existing VPC:

    https://pulsezta.blob.core.windows.net/gateway/templates/GCP/24-1-371/ivanti-zta-2-nics-existing-vpc.zip

    https://pulsezta.blob.core.windows.net/gateway/templates/GCP/24-1-371/ivanti-zta-3-nics-existing-vpc.zip

  - To deploy in a new VPC:

    https://pulsezta.blob.core.windows.net/gateway/templates/GCP/24-1-371/ivanti-zta-2-nics-new-vpc.zip

    https://pulsezta.blob.core.windows.net/gateway/templates/GCP/24-1-371/ivanti-zta-3-nics-new-vpc.zip

  > You can also choose to download Gateway templates through the **Gateways Overview** page of the *Controller* after you have defined the Gateway record. The opportunity to do this occurs later in this process.

- Credentials for the Google Cloud Platform Console.

  > These credentials must include sufficient permissions to create a virtual machine from a template image.

  Additionally, to manually specify GCP Gateway network interface settings:

- The primary (and optional secondary) DNS server IP address, and search domain.
- The required internal/private subnetworks must already be defined on Google Cloud Platform, including firewall settings. All required firewall settings for this interface are shown below.

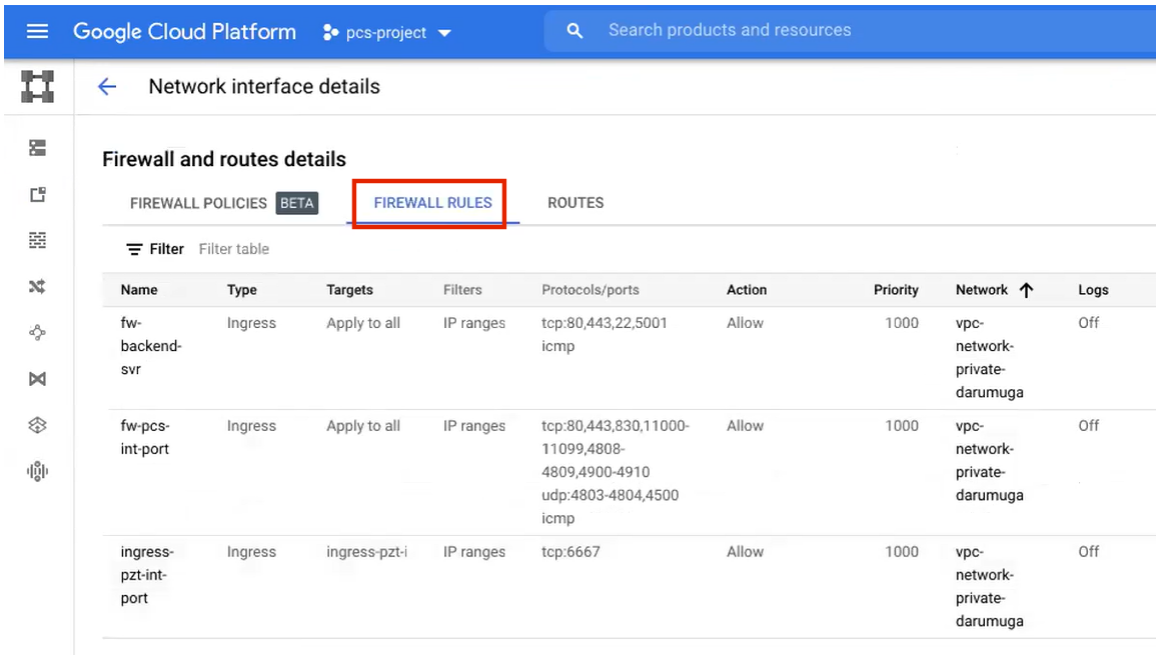

Internal/Private Firewall Rules

Refer to the Google Cloud Platform documentation for details.

- The required external/public subnetworks must already be defined on Google Cloud Platform, including firewall settings. All required firewall settings for this interface are shown below.

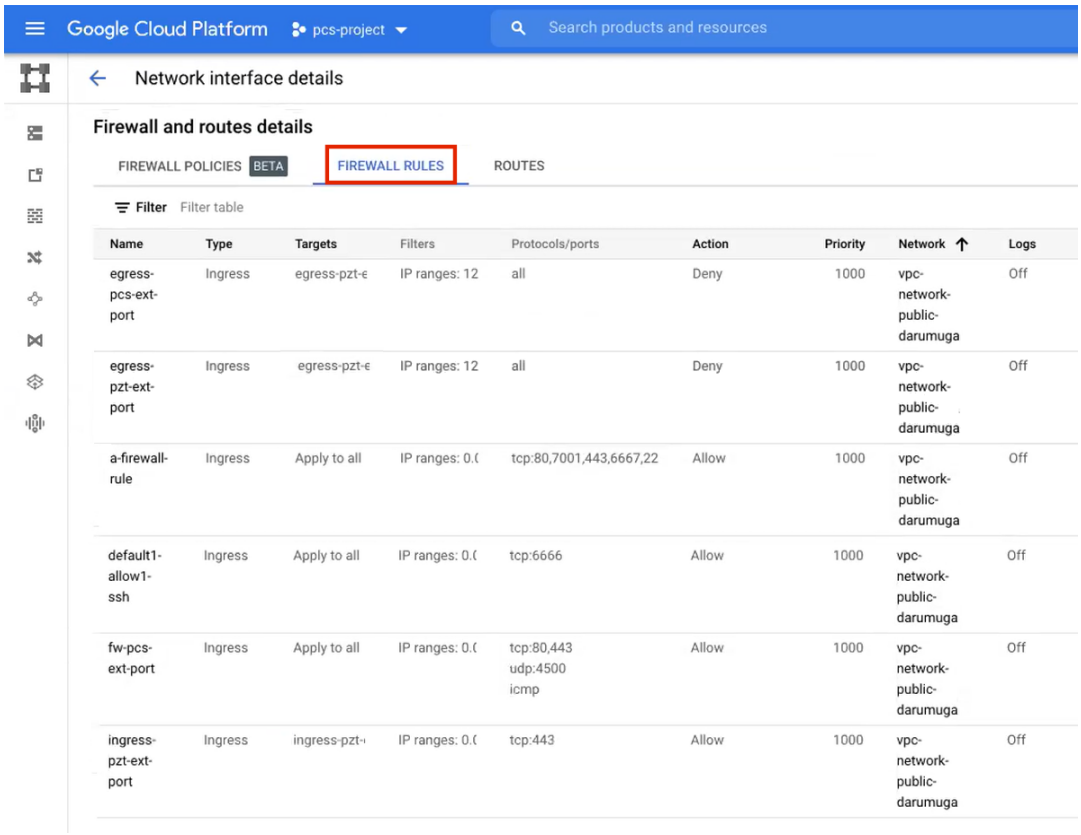

External/Public Firewall Rules

Refer to the Google Cloud Platform documentation for details.

- *(Optional)* Any required management subnetwork must already be defined on Google Cloud Platform, including firewall settings. All required firewall settings for this interface are shown below.

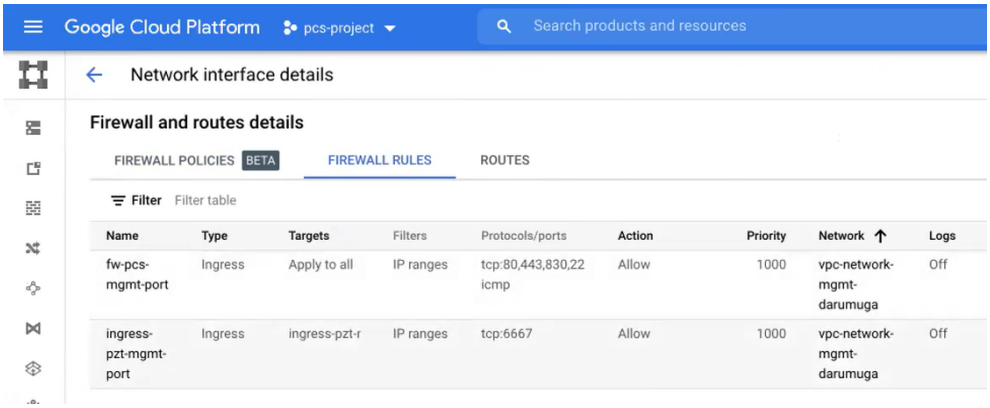

Management Firewall Rules

Refer to the Google Cloud Platform documentation for details.

After you have all required information, you can set up a *nZTA* GCP gateway, see "Adding a GCP Gateway" below.

Adding a GCP Gateway

To registering a Gateway on your *Controller*, use the **Gateway Details** dialog.

To begin, log into the *Controller* Tenant Admin Portal using the credentials provided in your welcome email. Two outcomes are possible:

- On unconfigured *nZTA* systems, the *Secure Access Setup* Onboarding wizard appears (see Working with the Onboarding Wizard). In this case, click **Add Gateway**.
- On configured *nZTA* systems, the *Network Overview* page appears. In this case:

1. From the nZTA menu, click the **Secure Access** icon, then select **Manage Gateways**.

   The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*

2. To add a new Gateway, select **Create** from the top-right:

   Add a new Gateway or Gateway Group

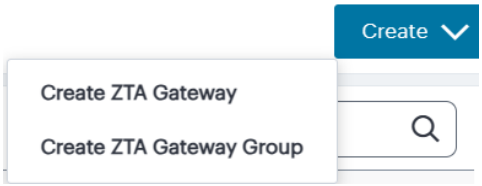

3. In the drop-down menu, click **Create ZTA Gateway**.

The **Gateway Details** dialog appears.

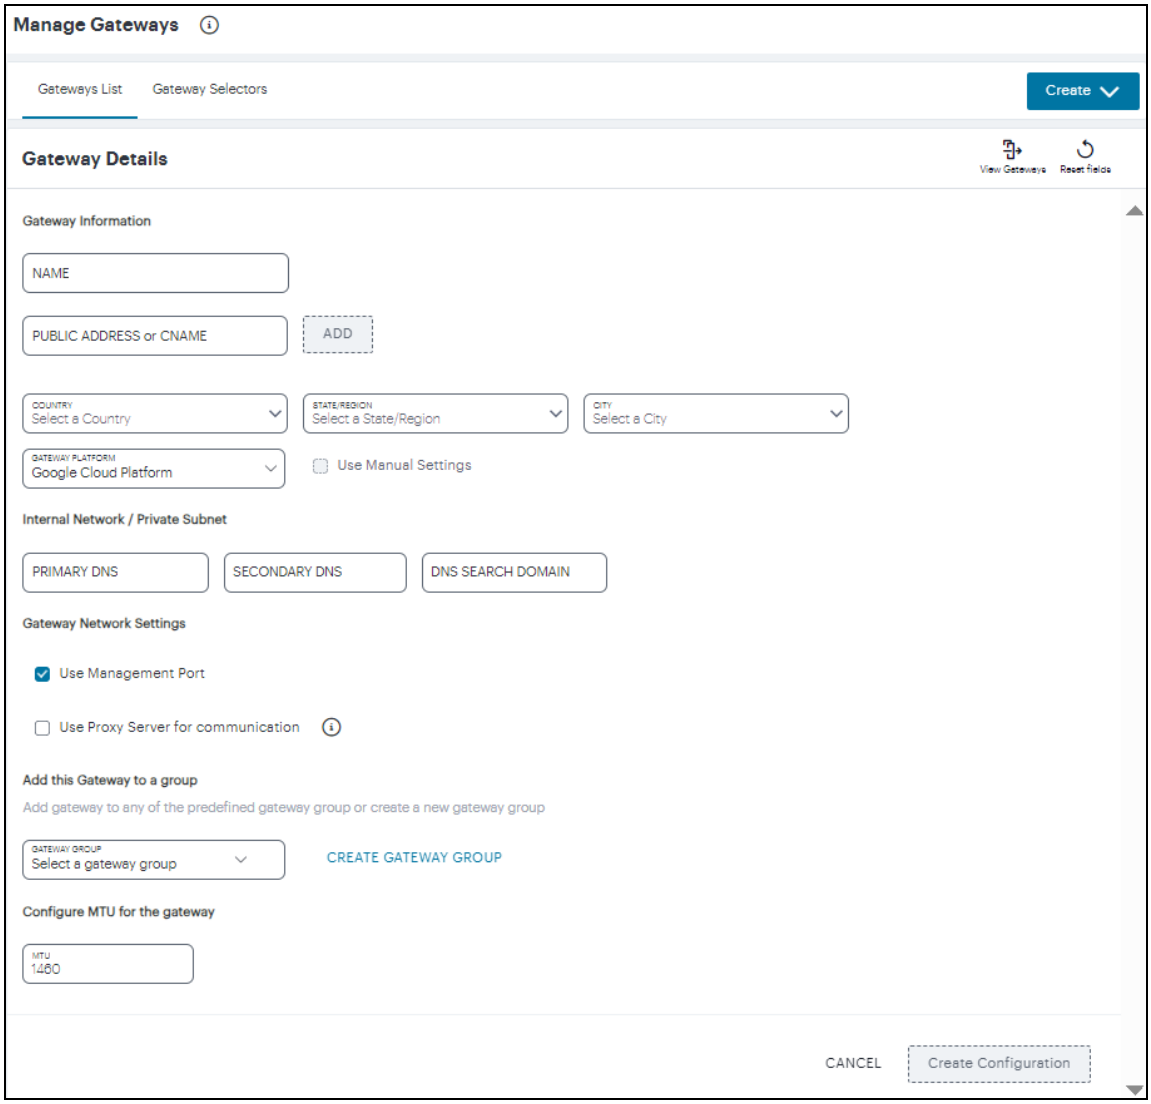

Gateway Details

Enter the following details:

1. Enter a **Name** for the Gateway.
2. Enter one or more **Public Address or CNAME** (Public IP address or CNAME) for the Gateway. Select **Add** to add each entry to the list. To learn more about this setting, see Configuring Networks in your Gateway Datacenter.

   If you want Google Cloud Platform to allocate a public IP address automatically, you can use a dummy IP address (for example, 1.1.1.1) at this point. You must then update the *Controller* with the allocated public IP address after the GCP VM instance is created, see Completing the Configuration of the Controller.

3. Select the geographic location details for the Gateway.
4. For **Gateway Platform**, select "Google Cloud Platform".
5. Enter the Primary DNS IP address for the Gateway.
6. (Optional) Enter the Secondary DNS IP address for the Gateway.
7. Enter the DNS Search Domain for the Gateway.
8. (Optional) Select the **Use Management Port** check box to use management network ports for *nZTA* traffic rather than internal ports.

   When the management port is enabled, Gateway will use management interface to communicate with *Controller* and NTP Server.
   The Gateway will still use the internal port for DNS resolution and NTP server name resolution.
   If the internal DNS cannot resolve the *Controller* domain, the internal interface will require internet access.

9. (Optional) Select the **Use proxy server for communication** check box to enable nZTA to *Controller* communication via proxy server. Proxy is supported on both internal and management interfaces of the Gateway. Once enabled, enter host name and port. Optionally, if your proxy server requires further authentication, enter a username and password to log in to the proxy server.

   Admin can configure proxy for existing Gateway after upgrading it to 22.4R3 version or later.

10. (Optional) Select a **Gateway Group** to which the new Gateway is to be added. To learn more about Gateway Groups, see Adding Gateway Groups for High Availability.

    A Gateway Group may have a defined public IP address, which you can specify as the **Public Address**.

11. **Configure MTU for the gateway**: Configurable MTU size allows admin to modify the default setting of nZTA gateways wherever it is needed. The value allowed is in the range of 576 to 1460 (IPv4).

To add a Gateway definition based on the settings you specified in this dialog, select **Create Configuration**.

You can now download your metadata, see Downloading Metadata for Google Cloud Platform.

In case of registration failure due to Gateway configuration mistakes in firewall rules, DNS, etc., you can re-register the gateway. It does not require re-deploying of Gateway. For details see "Re-registering a GCP Gateway" below.

Re-registering a GCP Gateway

Re-registration of Gateway includes the following capabilities:

- Gateway triggers re-registration on every launch of the gateway if the registration with controller fails or any update is made to the configuration parameters.
- The "View registration error report" option provides the reason for failure and solution to rectify it.

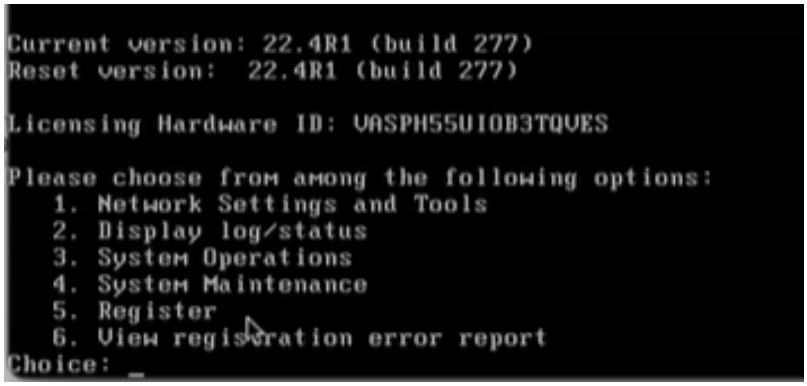

- On registration failures, admin is provided with the "Register" option to trigger the registration manually along with the existing debugging options such as networking tools, reboot, etc. This option can be used after rectifying any external issues such as network reach issue or Firewall rules following controller traffic from Gateway.
- To rectify registration failure due to the config error, first update the config settings in the Controller and download the config file. Then edit the Gateway instance and update the Custom metadata with the downloaded config file, and then reboot Gateway.

Sample screen: Update Config Value

An option is provided to regenerate and download the gateway init config from the controller admin interface.

Downloading Metadata for Google Cloud Platform

The preparation of metadata for use on Google Cloud Platform currently requires some manual steps:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways > Gateways List**.

   The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*.
3. Locate and select your GCP Gateway.

   The *Gateways Overview* page appears.

4. Click the **Download** icon, then choose **Download gateway init config** to obtain a copy of the Gateway definition file.

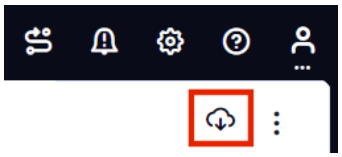

The Download Icon

5. Specify a save location for your Gateway definition file.

> The Gateway definition file is valid for 24 hours. If this period expires, you must replace the Gateway to generate a new Gateway definition file.

6. (Optional) If you have not yet downloaded the latest version of your Gateway VM image and optional YAML templates, click the **Download** icon and select **Download gateway VM image**. Save the archive file and unpack to a local workstation. Make sure the resulting file set is accessible from the Google Cloud Platform Management Portal.

You can now create a GCP gateway VM in Google Cloud Platform, see Uploading the GCP Virtual Machine Image onto the Google Cloud Platform.

Uploading the GCP Virtual Machine Image onto the Google Cloud Platform

To upload a GCP Gateway virtual machine image into Google Cloud Platform:

1. Access the *Google Cloud Platform Management Portal*, either from a client or a web browser, and log in using your Google Cloud Platform credentials.
2. In the Google Cloud Platform console, select your required project from the pull-down list on the title bar. For example:

GCP Select Project

3. Click the **Navigation** menu, and then select **Cloud Storage > Browser**.

   A list of GCP storage buckets appears.

4. Select the bucket into which you wish to place the GCP image.

   A page listing the current contents of the bucket appears.

5. (Optional) Navigate to the required folder within the bucket.

6. Click **Upload Files**. For example:

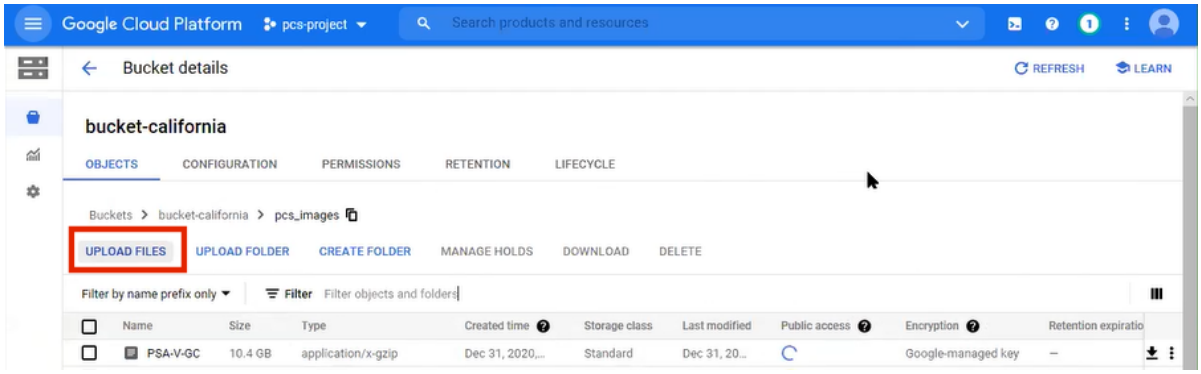

   GCP Upload Files

   An upload dialog appears.

7. Select the *ZTA Gateway* GCP virtual machine image *.tar* file from your local workstation (see Preparing to Create a GCP Gateway.), and click **Open**.

   > If you want to use the provided YAML templates to automate the creation of your VM instance (see Creating a VM Instance of the Uploaded GCP Image Using a Script/Template.), select these in addition to the image archive.

   The image archive and any selected template files are added to the bucket.

8. Wait until the upload completes. This may take several minutes.

9. Start a command line session from the title bar. For example:

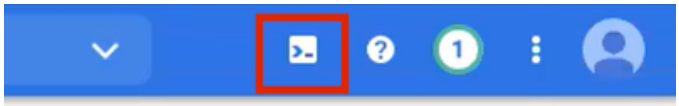

   GCP Upload Files

   A command line session starts.

10. Navigate to the project folder.

11. Create an image from the *nZTA Gateway* image archive using the following command:

```
gcloud compute images create <instance_name> --source-uri=gs://<bucket_
name>/<optional_path>/<image_name>.tar --guest-os-features MULTI_IP_SUBNET
```

You can now create a VM instance of the uploaded GCP image. To do this, either:

- Perform the task manually, see Creating a VM Instance of the Uploaded GCP Image Manually.
- Perform the task with a script/template, see Creating a VM Instance of the Uploaded GCP Image Using a Script/Template.

Creating a VM Instance of the Uploaded GCP Image Manually

This section describes how to manually create a virtual machine instance of the *nZTA Gateway* image inside Google Cloud Platform. You can also perform this process automatically using a script/template, see Creating a VM Instance of the Uploaded GCP Image Using a Script/Template.

1. Click the **Navigation** menu, and then select **Compute Engine > Images**.

   The **Images** page appears. For example:

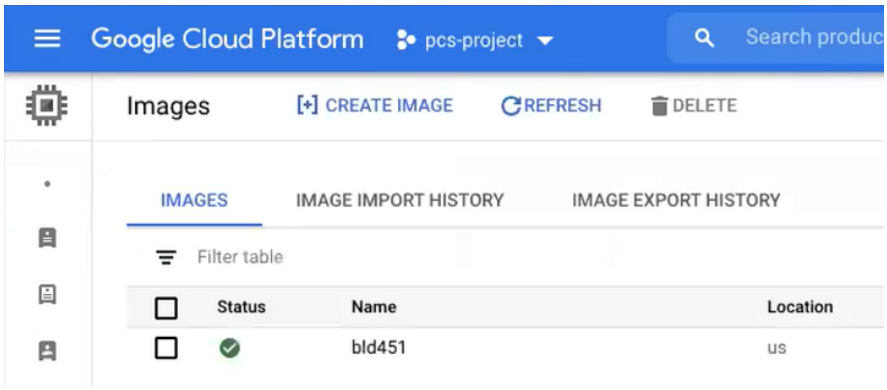

   GCP Images Page

2. Locate the new image in the list of images.

3. At the end of the image entry, click the action menu and select **Create Instance**.

   The **Create Instance** page appears. For example:

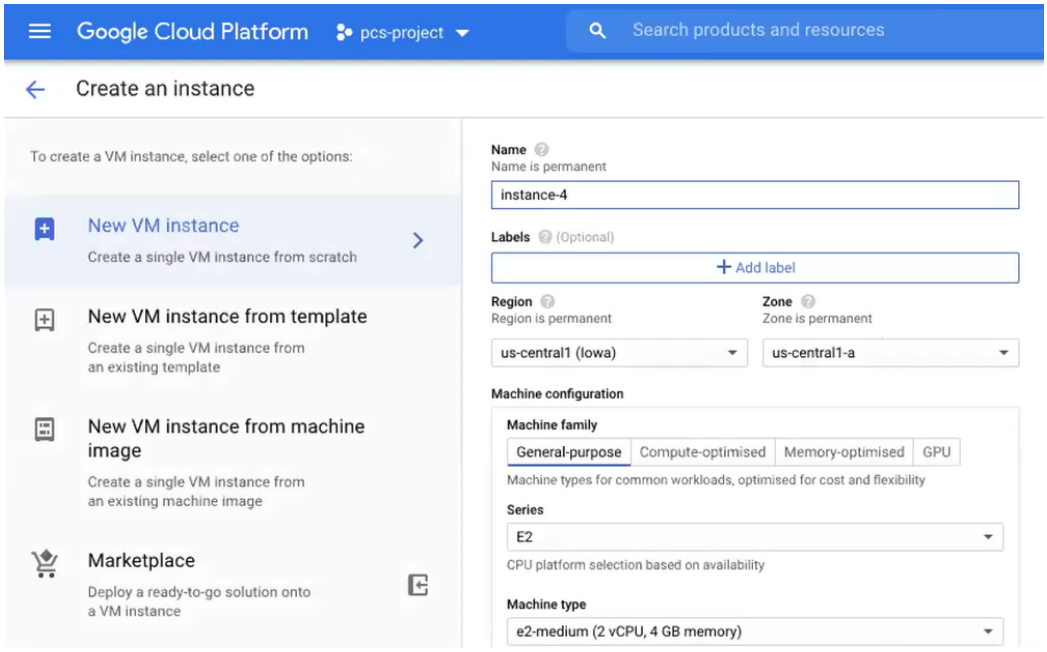

GCP Create Instance

4. On the **Create Instance** page:
    - Enter a **Name** for the new instance.
    - Select a **Region** and **Zone**.
    - Under **Machine configuration**:
        - For **Series**, select *N1*.
        - For **Machine Type**, select a minimum of *N1-Standard-n2* for 2-NIC, minimum of *N1-Standard-n4* for 3-NIC.
        - For **Boot Disk**, confirm that the correct image is already selected.
        - For **Firewall**, select the required HTTP/HTTPS options.
    - Expand the **Management, security, disks, networking, sole tenancy** options.
    - Select the **Management** tab.
    - Under **Metadata**:
        - For **Key**, enter *pulse-config*.
        - For **Value**, paste the text of the metadata file you downloaded earlier.
    - Select the **Networking** tab.
    - Under **Network interfaces**, click the **Edit** icon to change the default network interface selection.

      The **Network interface** options appear.
    - Under **Network interface**, specify a *private* (internal) network interface:
        - For **Network**, select the required private VPC.
        - For **Subnetwork**, select the required subnetwork.
        - Click **Done** to confirm the settings for the private network interface.
    - Under **Network interfaces**, click **Add network interface**.

      The **Network interface** options appear.

- Under **Network interface**, specify a *public* (external) network interface:
  - For **Network**, select the required public VPC.
  - For **Subnetwork**, select the required subnetwork.
  - Click **Done** to confirm the settings for the public network interface.
- (Optional) Click **Add network interface** and specify a management network interface.
- Click **Create** to confirm the settings and instantiate a VM instance of the image.

  The **VM Instances** page appears. This page shows the new VM instance of the image. For example:

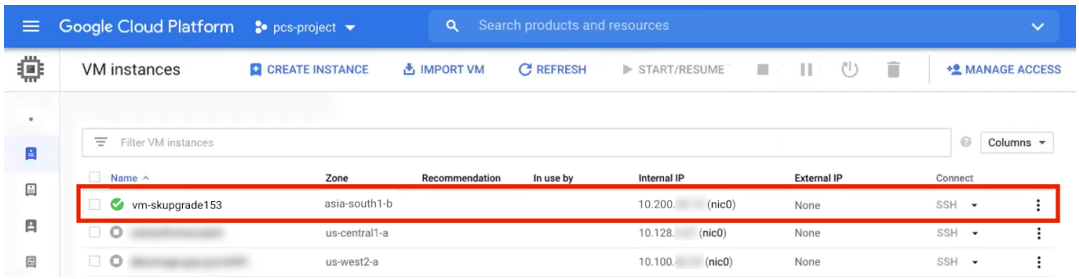

GCP Create Instance

5. On the **VM Instances** page, wait until the creation of the VM instance completes. This may take several minutes.
6. After the VM instance is created, click on it in the list of VM instances.

   The **VM instance details** page appears for the instance.
7. Confirm the details for the VM instance, including the number of network interfaces.
8. Make a note of the public IP address of the EXT interface (typically, this is *nic1*. This is required inside *nZTA*.

9. Under **Network interfaces**, confirm that the firewall settings from your VPCs are present for your specified network interfaces:

    - Click *nic0*. A summary page for this network interface appears.

      Under **Firewall and route details**, click the **Firewall Rules** tab and confirm that the following firewall rules are defined.

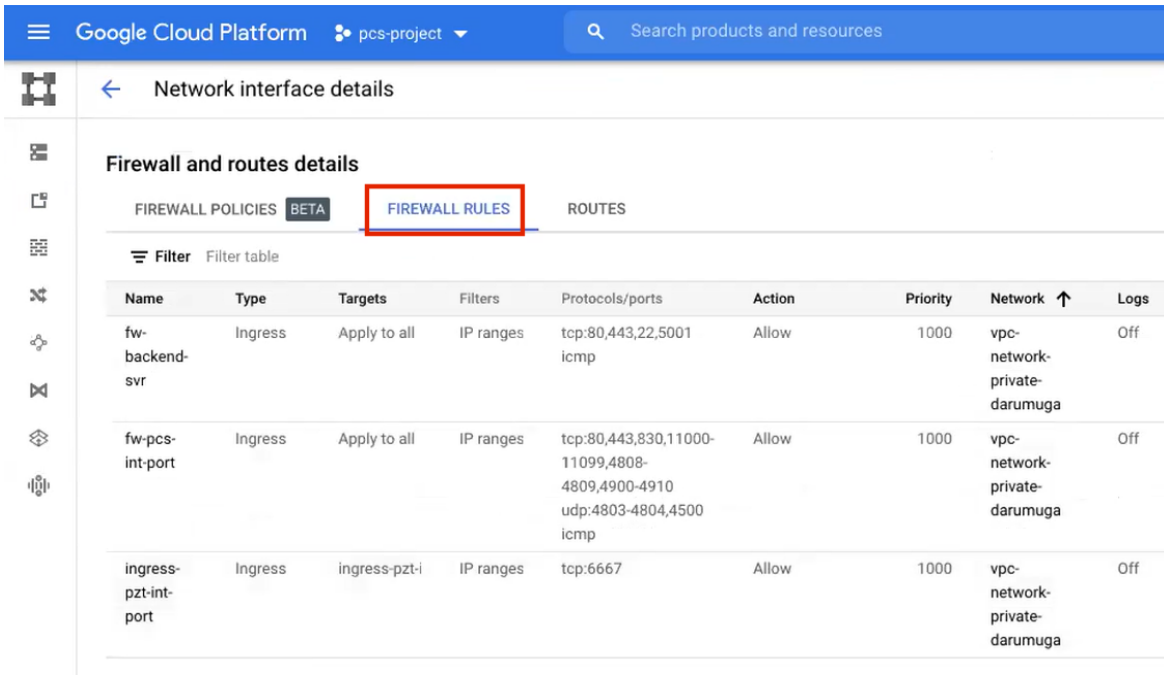

NIC0 Firewall Rules

- Click *nic1*. A summary page for this network interface appears.

  Under **Firewall and route details**, click the **Firewall Rules** tab and confirm that the following firewall rules are defined.

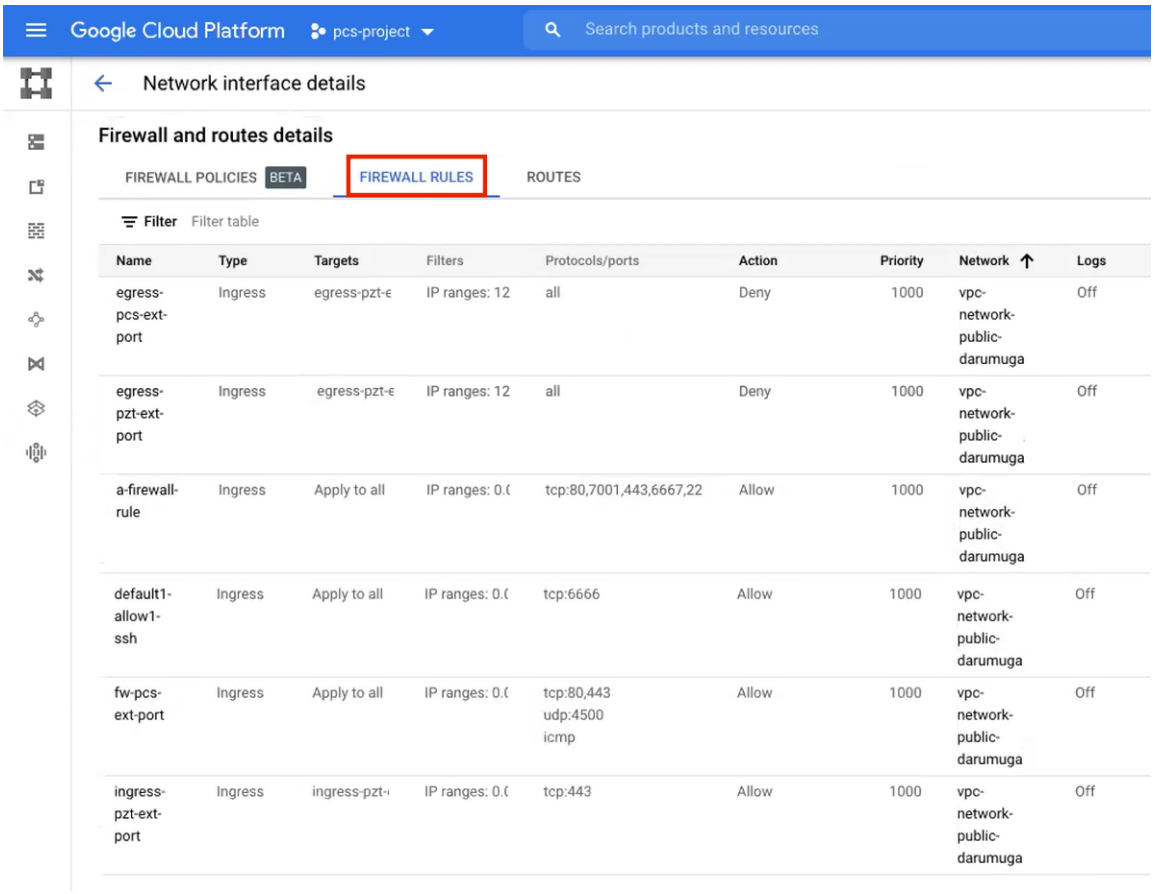

NIC1 Firewall Rules

- (Optional) Click *nic2*. A summary page for this optional network interface appears.

  Under **Firewall and route details**, click the **Firewall Rules** tab and confirm that the following firewall rules are defined.

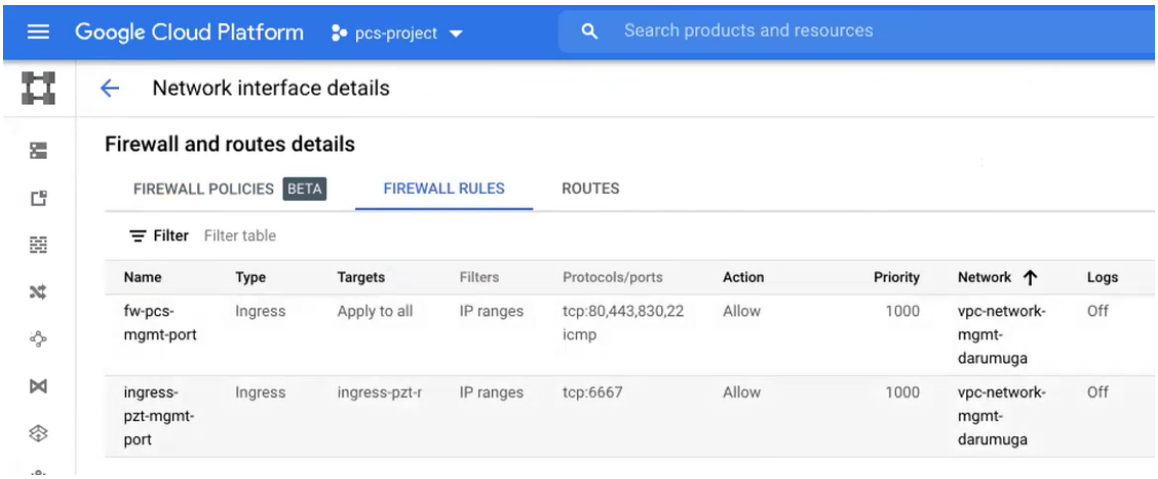

NIC2 Firewall Rules

10. The *VM instance details** page, click **Connect to serial console**

    A console monitor view (in a separate browser tab) shows the ongoing boot-up process for the instance.

11. Wait until the instance boot up is complete, and shows a screen similar to the following:

```
Welcome to the Pulse Zero Trust Access Serial Console!


Current version: 21.2R1 (build 153)
Rollback version: 21.2R1 (build 107)
Reset version:  21.2R1 (build 107)

Licensing Hardware ID: VASPH80EQ02HBPTES

Please choose from among the following options:
   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:
```

GCP Instance Console Monitor

You can then complete this process by updating the Gateway details on the *Controller*, see Completing the Configuration of the Controller.

Creating a VM Instance of the Uploaded GCP Image Using a Script/Template

This section describes how to automatically create a virtual machine instance of the *nZTA Gateway* image inside Google Cloud Platform using a script/template. You can also perform this process manually, see Creating a VM Instance of the Uploaded GCP Image Manually.

Ivanti provides YAML-based templates to create an instance of the *nZTA Gateway* image in the following configurations:

- Two network interfaces in an *existing* VPC.
- Three network interfaces in an *existing* VPC.

  Download:

  https://pulsezta.blob.core.windows.net/gateway/templates/GCP/24-1-371/ivanti-zta-2-nics-existing-vpc.zip

  https://pulsezta.blob.core.windows.net/gateway/templates/GCP/24-1-371/ivanti-zta-3-nics-existing-vpc.zip
- Two network interfaces in a *new* VPC.
- Three network interfaces in a *new* VPC.

  Download:

  https://pulsezta.blob.core.windows.net/gateway/templates/GCP/24-1-371/ivanti-zta-2-nics-new-vpc.zip

  https://pulsezta.blob.core.windows.net/gateway/templates/GCP/24-1-371/ivanti-zta-3-nics-new-vpc.zip

  You can obtain these templates through the links given here, or as part of the archive file set provided through the *Download* link on the **Gateways Overview** page in the *nZTA* Tenant Admin Portal after you have defined a Gateway record.

To use a template:

1. Download the required template archive file to your local workstation.
2. Unpack the downloaded archive file to a location that is accessible from Google Cloud Platform. Each archive contains three files. For example, for the two-interface (existing VPC) version of the archive:

```
 pulsesecure-zta-2nics-existing-vpc.jinja
pulsesecure-zta-2nics-existing-vpc.jinja.scheme
pulsesecure-zta-2nics-existing-vpc.yaml
```

3. Edit the YAML file `properties` section to reflect your project and instance requirements, including the `user_data` property.

   An example of an *existing VPC* YAML file is provided here:

```
          imports:
  - path: pulsesecure-zta-2-nics-existing-vpc.jinja
resources:
  - name: my-vm
    properties:
      project: zta-gw-263035
      email: admin@example.com
      region: asia-south1
      zone: asia-south1-b
      image: ztagcp123
      machine_type: n1-standard-2
      int_network:
      ext_network:
      int_subnetwork:
      ext_subnetwork:
      user_data:
    type: pulsesecure-zta-2-nics-existing-vpc.jinja
```

   An example of a *new VPC* YAML file is provided here:

```
          imports:
  - path: pulsesecure-zta-2-nics-new-vpc.jinja
resources:
  - name: my-vm
    properties:
      deploy_with_lb: yes
      project: zta-gw-263035
      email: admin@example.com
      region: asia-south1
      zone: asia-south1-b
      image: ztagcp123
      machine_type: n1-standard-2
      user_data: <pulse-config><primary-dns>8.8.8.8<\primary-dns> ...
```

```
    int_cidr: 192.0.2.0/24
    ext_cidr: 192.0.2.0/24
  type: pulsesecure-zta-2-nics-new-vpc.jinja
```

Where you are specifying a new VPC for your virtual machine instance, make sure you use properties (for example, networking settings) that do not conflict with an existing VPC.

The following table lists all possible template properties and their meaning:

Property	Description	ExampleValue
project	Project Identifier	myproject
email	Registered service account email address	email@example.com
region	Thename of the region in which you want to deploy your VM instance	asia-east1
zone	The name of the zone in which you want to deploy your VM instance	asia-east1-b
image	Virtual machine image name	gwimage
machine_type	GCP machine type	For 2-nic: N1-Standard-n2 minimum For 3-nic: N1-Standard-n4 minimum
int_network	VPC network name for internal network	nw1-private
ext_network	VPC network name for external network	nw1-public
mgmt_network	VPC network name for management network	nw1-mgmt
int_subnetwork	Subnet name for internal VPC	int-nw1

Property	Description	ExampleValue
ext_subnetwork	Subnet name for external VPC	ext-nw1
mgmt_subnetwork	Subnet name for management VPC	mgmt-nw1
user_data	The Gateway config file downloaded from the *nZTA* controller	<pulse-config>... </pulse-config>
int_cidr	IP address range for the internal subnet	192.0.2.0/24
ext_cidr	IP address range for the external subnet	198.51.100.0/24
mgmt_cidr	IP address range for the management subnet	203.0.113.0/24
health_check	Health check rule name with port 443	hc1
ingress_pzt_int_port	Firewall rule for inbound internal network	ing-int
ingress_pzt_ext_port	Firewall rule for inbound external network	ing-ext
egress_pzt_ext_port	Firewall rule for outbound external network	egress
ingress_pzt_mgmt_port	Firewall rule for inbound management network	ing-mgmt
vm_name	Virtual machine instance Name	vm77
router_int	Cloud router name for the internal interface	rtr1
nat_int	NAT Gateway name for the internal interface	nat1
router_mgmt	Cloud router name for the management interface	rtr-mgmt1

Property	Description	ExampleValue
nat_mgmt	NAT Gateway name for the management interface	nat-mgmt
ip_address	Load balancer front-end IP address	3nic
instance_group	Instance group name	group1
load_balancer	Load balancer name	lb
max_connections	The number of maximum connections	2
target_proxy	Targeted proxy name	proxy1
front_end	Load balancer front-end profile name	front1

4. Save the YAML file.
5. On the Google Cloud Platform, start a command line session from the title bar. For example:

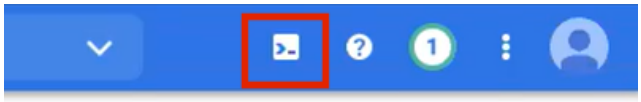

GCP Upload Files

A command line session starts.

6. Select the required project:

```
gcloud config set project <project-name>
```

7. Within the project folder, create a *deploymentmanager* folder.
8. Copy the three script files to this folder.

9. Create a VM instance from the *nZTA Gateway* image archive file using the following command:

   ```
   gcloud deployment-manager deployments create <vm-name> --config
   <yaml_file>
   ```

   For example:

   ```
   gcloud deployment-manager deployments create vm-gcp-123 --config
   pulsesecure-zta-3-nics-existing-vpc.yaml
   ```

10. Wait until the command completes.
11. On the **VM Instances** page, click on the new VM in the list of VM instances.

    The **VM instance details** page appears for the instance.

12. Confirm the details for the VM instance, including the number of network interfaces.
13. Make a note of the public IP address of the INT interface (typically, this is *nic0*. This is required inside *nZTA*.

You can now complete this process by updating the Gateway details on the *Controller*, see Completing the Configuration of the Controller.

Completing the Configuration of the *Controller*

If you specified a dummy public IP address (for example, *1.1.1.1*) when you created the Gateway on the *Controller*, you now need to update the *Controller* with the allocated public IP address for the Gateway VM instance on Google Cloud Platform.

> You do not need to perform the following procedure if you specified the correct public IP address when you created the Gateway on the *Controller*, see "Adding a GCP Gateway" on page 307.

1. Return to the **Gateways List** page in the *nZTA* Tenant Admin Portal.
2. Locate the new Gateway record in the list and confirm that its status has updated to *Connected*.
3. Click the Gateway and then select **Secure Access > Manage Gateways > *Gateway* > Configuration**.
4. Under **Gateway Network Settings**, delete the current public IP setting and replace it with the public IP address if the nic1 (external) interface for the VM instance.

5. (Optional) After you have registered a Gateway, you can configure it (or the Gateway Group to which it belongs) as the default Gateway if required. A default Gateway handles all requests from applications that are not referenced by any secure access policy. See Configuring a Default Gateway for Application Discovery.

Workflow: Creating a Gateway in Oracle Cloud Platform

This workflow leads you through the processes for setting up a Gateway on the Oracle Cloud Platform (OCI).

These processes must be performed in sequence:

- Preparing to create an Oracle gateway, see "Preparing to Create an Oracle Gateway" below.
- Creating the gateway record in the Controller, see "Adding an Oracle Gateway" on page 330.
- Downloading Metadata for Oracle Cloud Platform, see "Downloading Metadata for Oracle Cloud Platform" on page 335.
- Uploading the Oracle Image onto the Oracle Cloud Platform, see "Uploading the Oracle Virtual Machine Image onto the Oracle Cloud Platform" on page 336.
- Creating a VM Instance of the OCI image:
    - Creating a VM Instance of the Uploaded Oracle Image Using a Script/Template, see "Creating a VM Instance of the Uploaded OCI Image Using Terraform Script" on page 340.
    - Refer Terraform Configurations for details on the configuration, see "Terraform Configurations" on page 342.
- Creating a VM Instance of the Uploaded OCI Image Using any Other Methods, see "Creating Gateway Selectors" on page 239

Preparing to Create an Oracle Gateway

Before you start, make sure you have the following information and files:

- An identifying name for the Gateway.
- The public IP address for the Gateway. This is the IP address at which clients can externally reach the Gateway instance, such as an LB/NAT or Datacenter network forward rules.

  If you want Oracle Cloud platform to allocate a public IP address automatically, you can use a dummy IP address (for example, 1.1.1.1) when you create the Gateway on nZTA.
- The Gateway geographic location.
- (Optional) The name of the Gateway Group to which you want to add this new Gateway record. To learn more about Gateway Groups, see "Adding Gateway Groups for High Availability" on page 238.

  Gateway Group may have a defined public IP address, which you can specify during the creation of the Gateway.
- The nZTA Gateway Oracle virtual machine image: https://pulsezta.blob.core.windows.net/gateway/ISA-V-OCI-ZTA-22.7R1-371.1.tar.gz

  Download a copy of the Oracle Gateway image as a compressed zip archive file, then decompress the archive to a local workstation. Make sure that the resulting file set is accessible from the Oracle Cloud Platform Console.

  You can also choose to download the Gateway image through the Gateways Overview page of the Controller after you have defined the Gateway record. The opportunity to do this occurs later in this process.
- (Optional) Oracle Gateway deployment scripts, suitable for automating the creation of your Oracle VM instances.

  Template files: https://pulsezta.blob.core.windows.net/gateway/templates/OCI/24-1-371/Terraform.zip
- Credentials for the Oracle Cloud Platform Console.

  These credentials must include sufficient permissions to create a virtual machine using terraform scripts.
- The primary (and optional secondary) DNS server IP address, and search domain.

Adding an Oracle Gateway

To register a Gateway on your Controller, use the Gateway Details dialog. To begin, log into the Controller Tenant Admin Portal using the credentials provided in your welcome email. Two outcomes are possible:

- On unconfigured nZTA systems, the Secure Access Setup Onboarding wizard appears (see "Working with the Onboarding Wizard" on page 56). In this case, click **Add Gateway**.

- On configured nZTA systems, the Network Overview page appears. In this case:

    1. From the nZTA menu, click the Secure Access icon, then select **Manage Gateways**.

       The Gateways List page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the Controller.

    2. To add a new Gateway, select **Create** from the top-right.

3. From the drop-down menu, click **Create ZTA Gateway**.

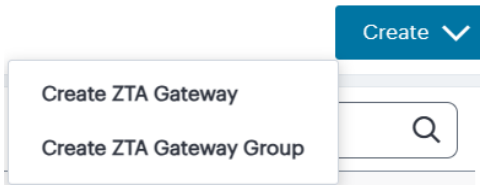

The Gateway Details dialog appears.

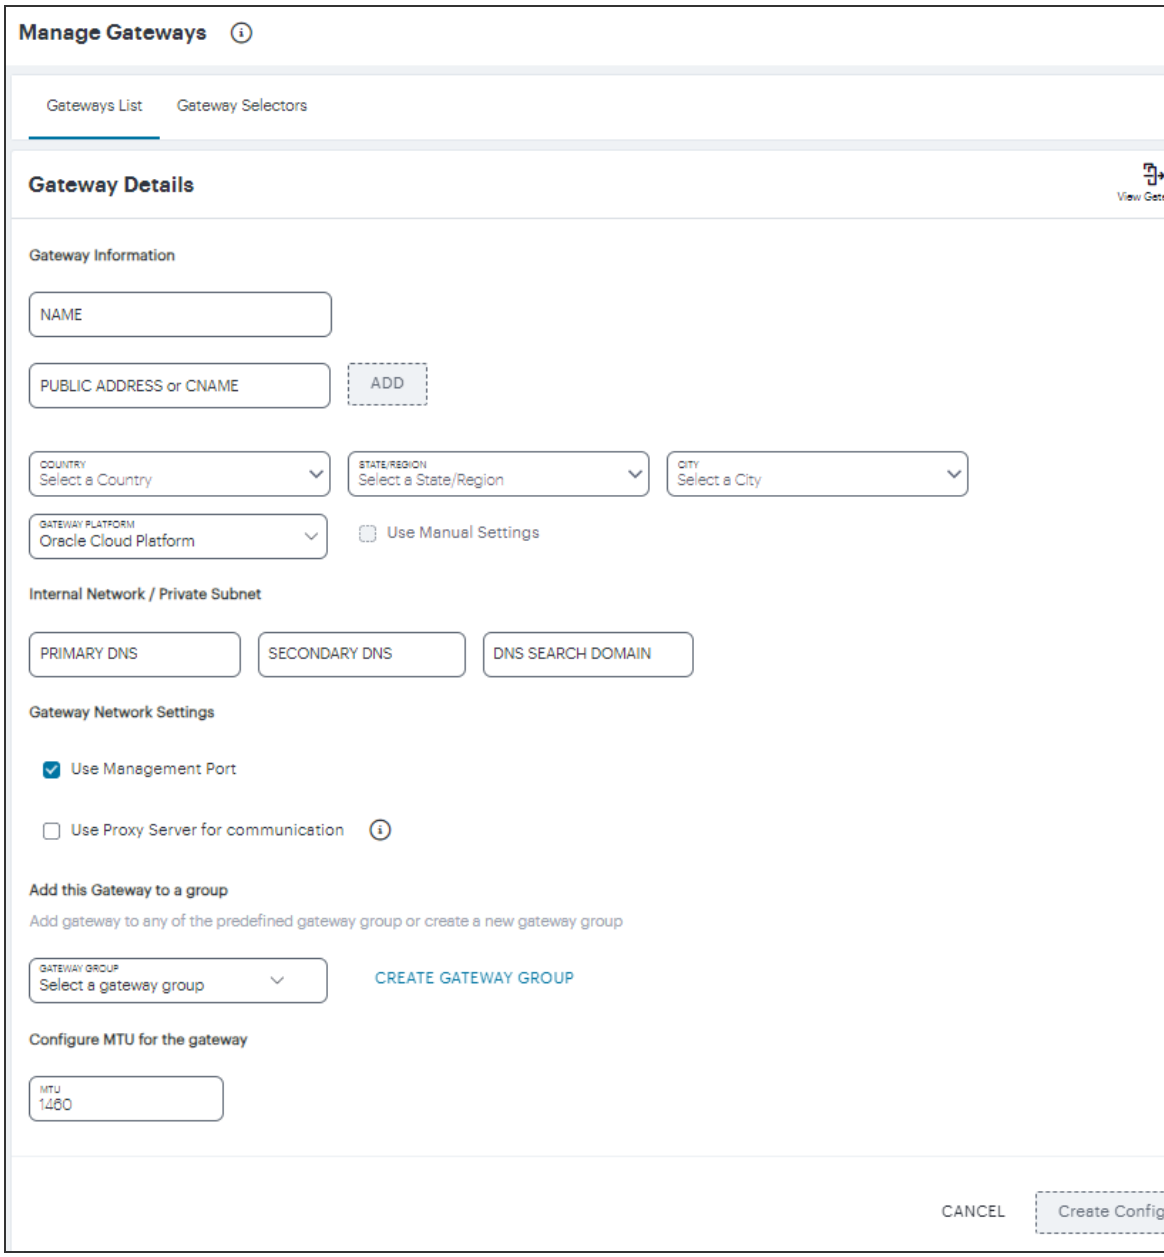

Enter the following details:

1. Enter a **Name** for the Gateway.
2. Enter one or more **Public Address or CNAME** (Public IP address or CNAME) for the Gateway.

3. Select **Add** to add each entry to the list. To learn more about this setting, see "Configuring Networks in your Gateway Datacenter" on page 214.

   If you want Oracle Cloud Platform to allocate a public IP address automatically, you can use a dummy IP address (for example, 1.1.1.1) at this point. You must then update the Controller with the allocated public IP address after the Oracle VM instance is created (if it is not updated automatically).

4. Select the geographic location details for the Gateway.
5. For Gateway Platform, select **Oracle Cloud Platform**.
6. Enter the Primary DNS IP address for the Gateway.
7. (Optional) Enter the Secondary DNS IP address for the Gateway.
8. Enter the DNS Search Domain for the Gateway.
9. (Optional) Select the **Use Management Port** check box to use management network ports for nZTA traffic rather than internal ports.

   When the management port is enabled, Gateway will use management interface to communicate with Controller and NTP Server.

   The Gateway will still use the internal port for DNS resolution and NTP server name resolution.

   If the internal DNS cannot resolve the Controller domain, the internal interface will require internet access.

10. (Optional) Select the **Use proxy server for communication** check box to enable nZTA to Controller communication via proxy server.

    Proxy is supported on both internal and management interfaces of the Gateway. Once enabled, enter host name and port.

    Optionally, if your proxy server requires further authentication, enter a username and password to log in to the proxy server.

11. (Optional) Select a Gateway Group to which the new Gateway is to be added. To learn more about Gateway Groups, see "Adding Gateway Groups for High Availability" on page 238.

    A Gateway Group may have a defined public IP address, which you can specify as the Public Address.

12. **Configure MTU for the gateway**: Configurable MTU size allows admin to modify the default setting of nZTA gateways wherever it is needed. The value allowed is in the range of 576 to 1500 (IPv4).

13. To add a Gateway definition based on the settings you specified in this dialog, select **Create Configuration**.

You can now download your metadata, see "Downloading Metadata for Oracle Cloud Platform" below.

In case of registration failure due to Gateway configuration mistakes in firewall rules, DNS, etc., you can re-register the gateway. It does not require re-deploying of Gateway.

Downloading Metadata for Oracle Cloud Platform

The preparation of metadata for use on Oracle Cloud Platform currently requires some manual steps:

1. Log into the Controller as a Tenant Admin, see Logging in as a Tenant Administrator.

2. From the nZTA menu, click the Secure Access icon, then select **Manage Gateways > Gateways List**.

   The Gateways List page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the Controller.

3. Locate and select your Oracle cloud Gateway.

   The Gateways Overview page appears.

4. Click the Download icon, then choose **Download gateway init config** to obtain a copy of the Gateway definition file.

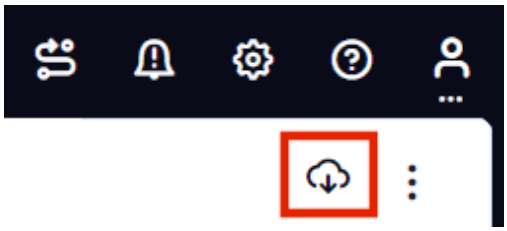

5. Specify a save location for your Gateway definition file.

   The Gateway definition file is valid for 24 hours. If this period expires, you must replace the Gateway to generate a new Gateway definition file.

6. (Optional) If you have not yet downloaded the latest version of your Gateway VM image and optional YAML templates, click the Download icon and select Download gateway VM image. Save the archive file and unpack to a local workstation. Make sure the resulting file set is accessible from the Oracle Cloud Platform Management Portal.

You can now create an Oracle cloud gateway VM in Oracle Cloud Platform, see "Uploading the Oracle Virtual Machine Image onto the Oracle Cloud Platform" below.

Uploading the Oracle Virtual Machine Image onto the Oracle Cloud Platform

To upload a Oracle Gateway virtual machine image into Oracle Cloud Platform:

1. Access the Oracle Cloud Platform Management Portal, either from a client or a web browser, and log in using your Oracle Cloud Platform credentials.
2. Click the Navigation menu, and then select **Storage** > **Buckets**.

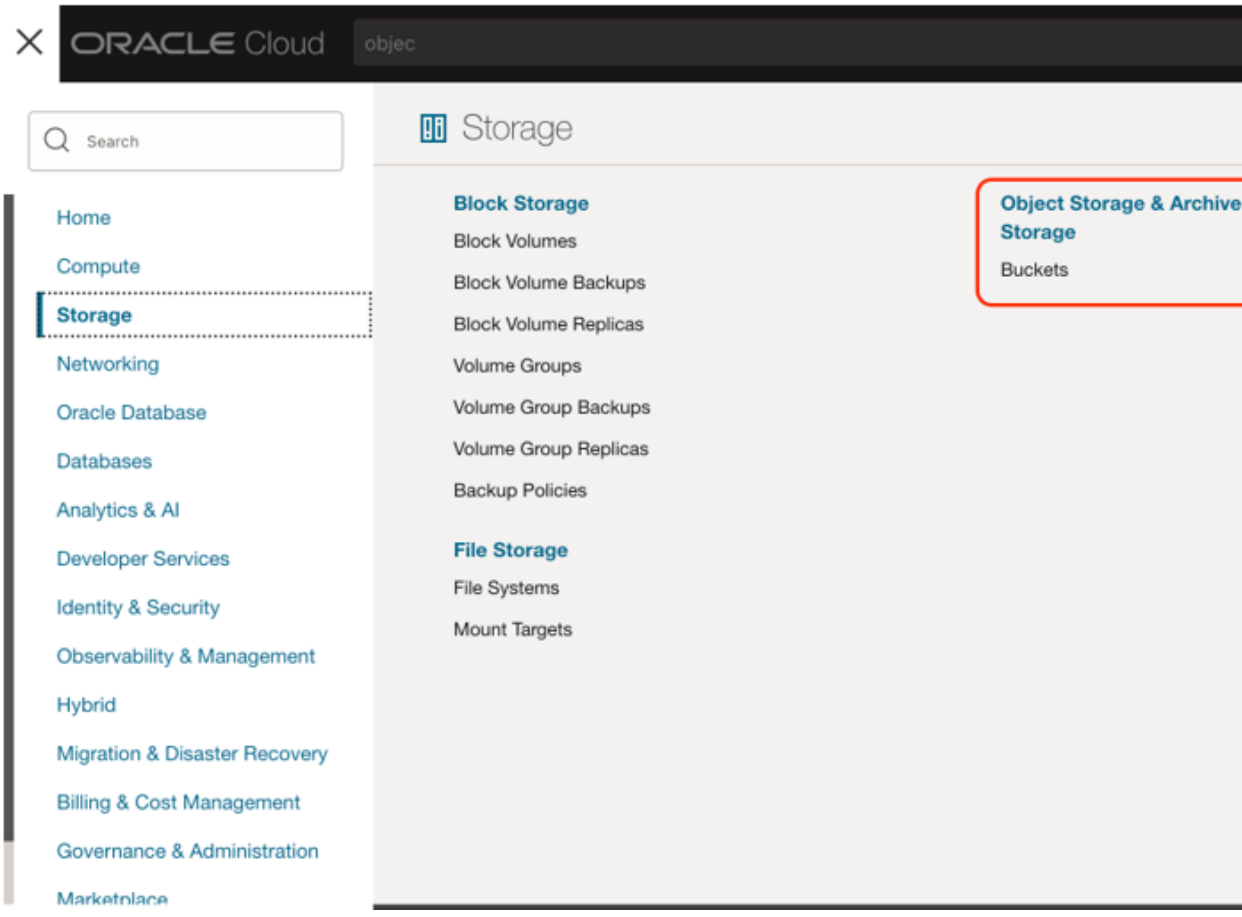

3. Select the right compartment and then the list of OCI storage buckets appears as below.
4. Select the bucket into which you wish to place the OCI image.

   A page listing the current list of the files from the bucket appears.
5. Click **Upload**.

   An upload dialog appears.

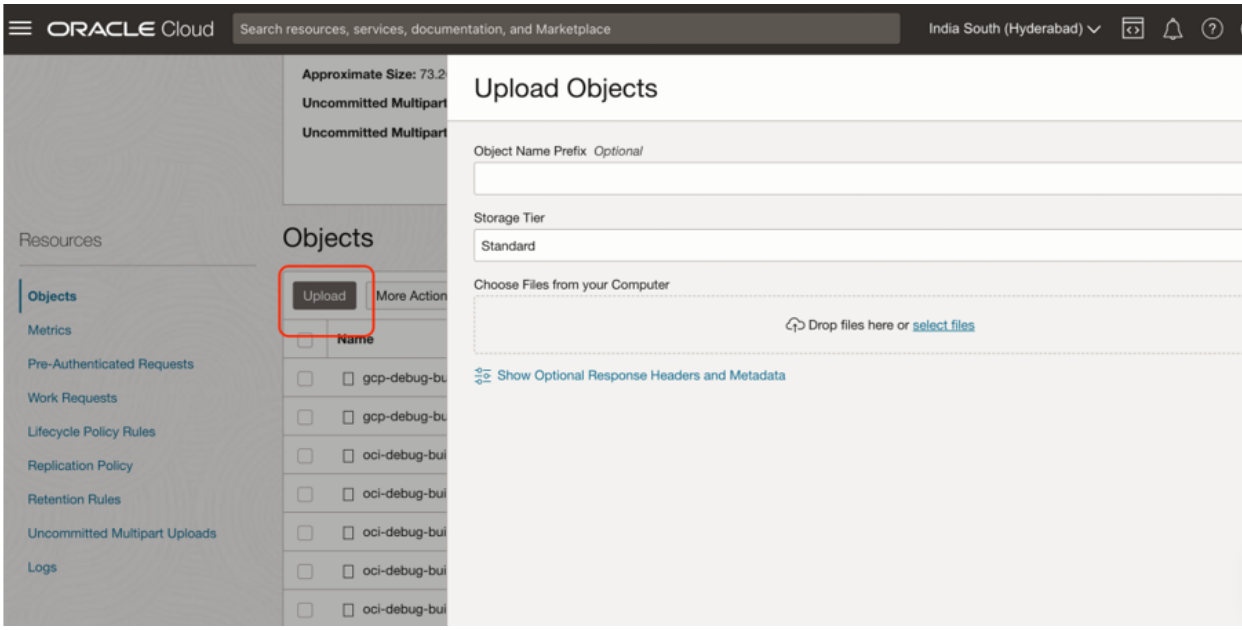

6. Select the *ZTA Gateway* OCI virtual machine image.tar file from your local workstation and click **Open**.
7. Wait until the upload completes. This may take several minutes.
8. Once upload completes. Create an image from the bucket using the following method:

1. Click the Navigation menu, and then select **Compute** > **Custom Images**.

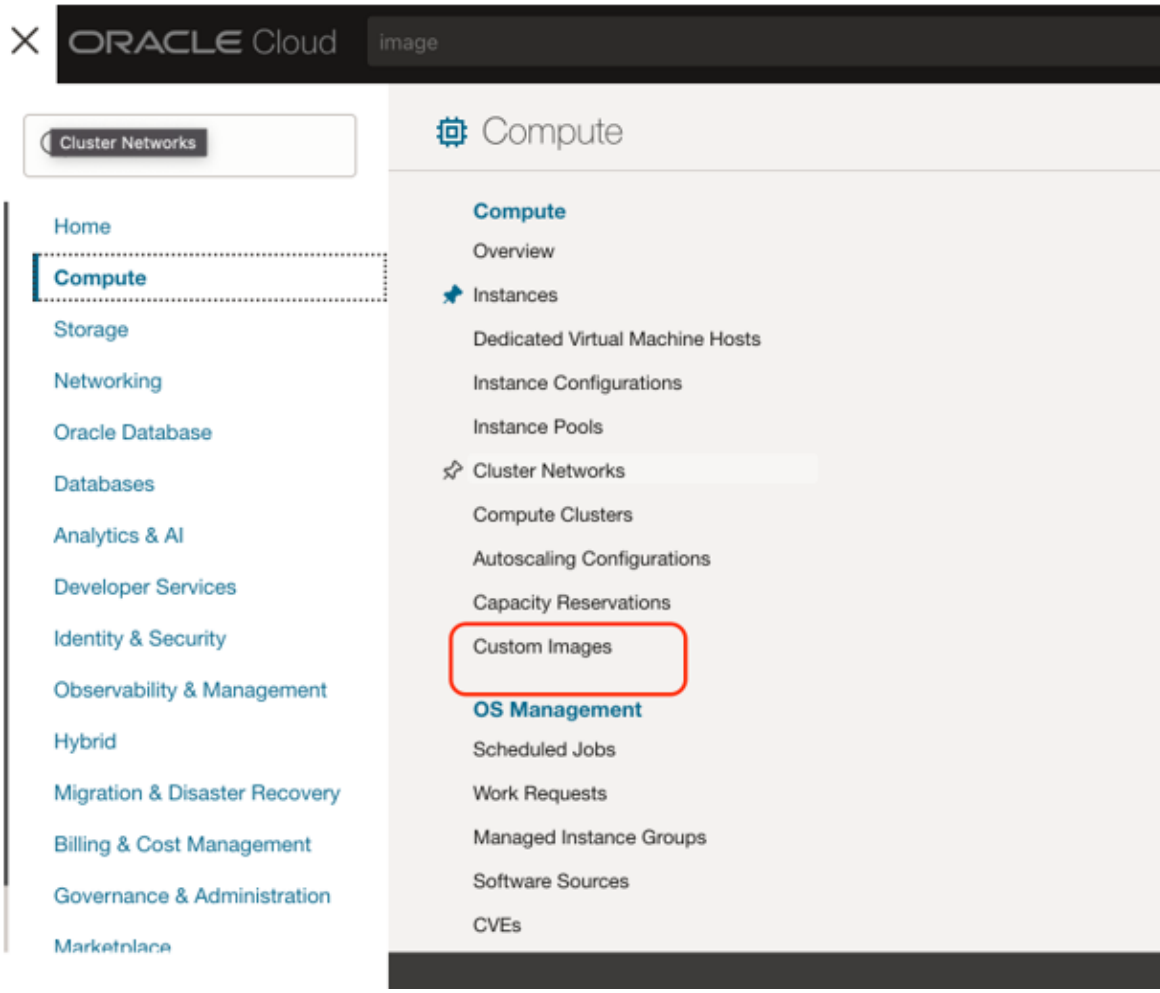

2. Select the right compartment and then the list of current images appears.

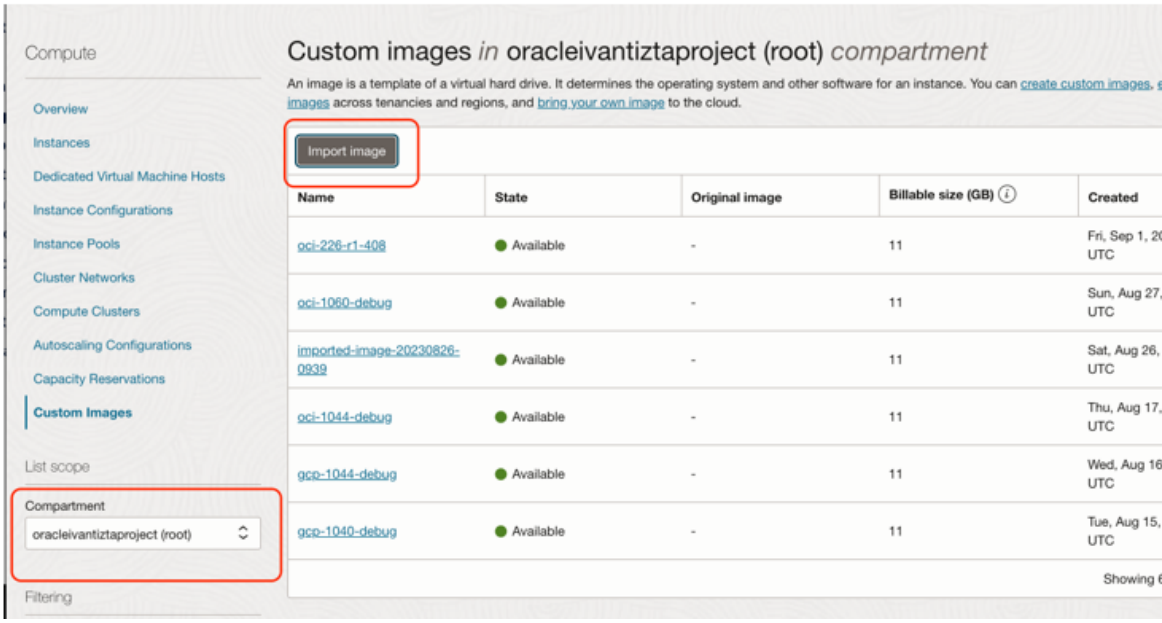

3. Click on **Import image**. An import dialog will appear and then choose the .tar file that was uploaded in the bucket. Refer to the below screenshots.

    - Ensure the OS is selected as CentoS.
    - Ensure Launch mode is chosen as Paravirtualized mode.
    - Ensure Image type is chosen as QCOW2.

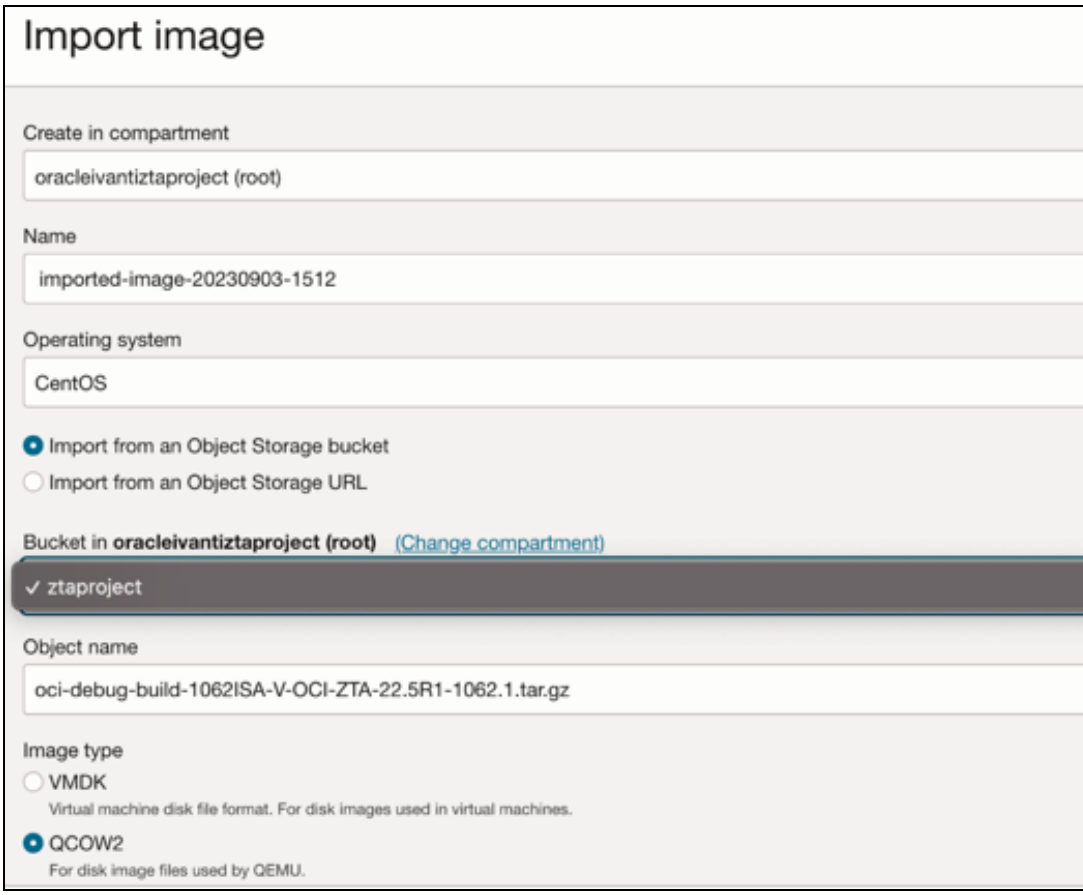

- Wait until the import completes. This may take several minutes.

You can now create a VM instance of the uploaded OCI image. To do this, either:

- Perform the task manually, see "Working with Gateways" on page 211.
- Perform the task with a script, see "Creating a VM Instance of the Uploaded OCI Image Using Terraform Script" below.

Creating a VM Instance of the Uploaded OCI Image Using Terraform Script

Pre-requisites: Ensure OCI configurations required for CLI access is enabled. This is one-time process. Please see here for details:

1. Download the required template archive file to your local workstation.
2. Upload all the scripts to Oracle cloud shell as below:

1. Open the cloud shell using the option as shown below.

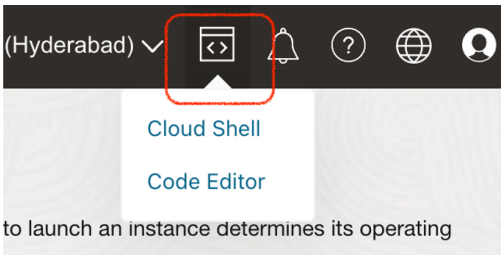

2. Upload the terraform scripts using the upload option from the cloud shell settings as shown below. Once it is uploaded, the scripts will be present in the user's home directory.

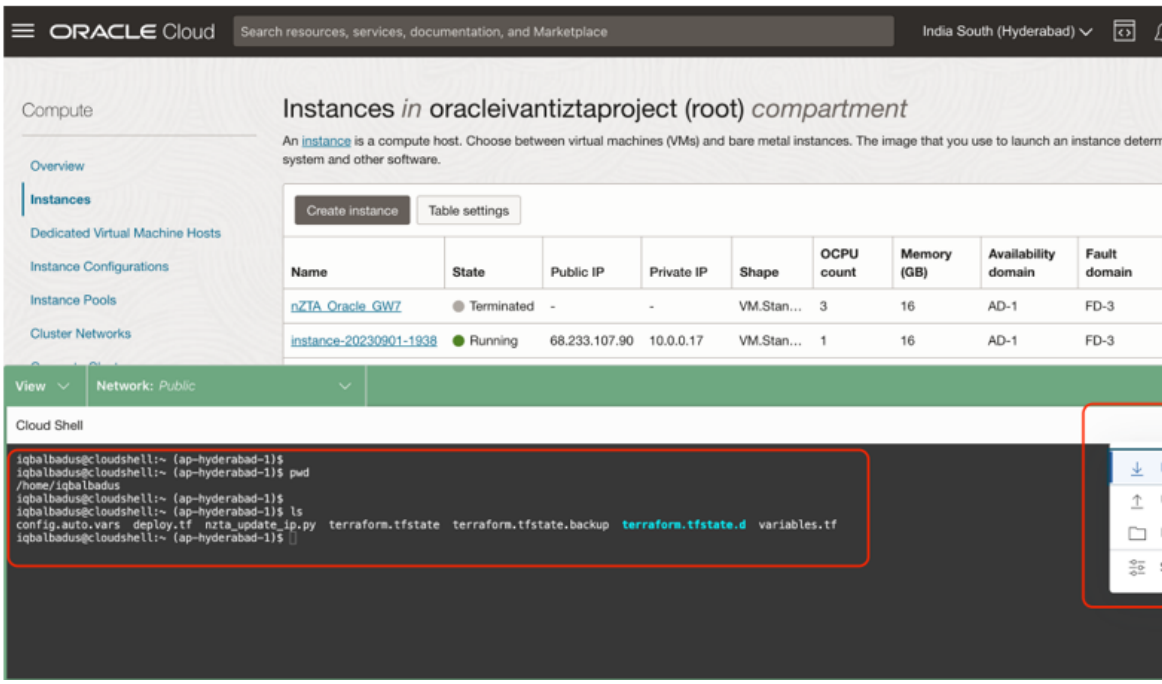

3. Edit the *config.auto.vars* file as per the intended deployment. Refer "Terraform Configurations" on the next page for details on the configuration.

3. Run **terraform init**.
4. Run **terraform validate**. This will let the admin know if there are any issues with the configurations.
5. Run **terraform apply**. This will trigger the deployment process.
6. If any resource deployment fails, retry running the command **Terraform apply** again.

Terraform Configurations

Configuration	Description	Sample Value
enable_management	This defines, if nZT	enable_management = true

Description	Sample Value
Needs an additional	

Configuration	Description	Sample Value
	nterface.	
2internal_exist	This de	internal_existing_vcn = false

Configuration	Description	Sample Value
ing_vcn	fines, if existin	

			Description	**Sample Value**
			obe used for inte	

			Description	**Sample Value**
			e, then we need to pr	

Configuration	Description	Sample Value
external_existing_vcn	This defines, if	external_existing_vcn = false

			Description	**Sample Value**
			existing VCN to b	

	Configuration	Description	Sample Value
		nal interface. If tru	

			Description	**Sample Value**
			iguration_external_vcn_	

Configuration	Description	Sample Value
management_existing_vcn	This is defines, if	management_existing_vcn = false

Configuration	Description	Sample Value
n	existing VCN to b	

Configuration	Description	Sample Value
	ement interface.ifIt	

Configuration	Description	Sample Value
	nfiguration management_	

Configuration	Description	Sample Value
use_internal_vcn_for_all	This defines, if in	use_internal_vcn_for_all = false

Configuration	Description	Sample Value
l	terminal VCN to be used	

Description	Sample Value
external/management).	

Configuration	Description	Sample Value
6internal_existing_subnet	This defines, if existing subn	internal_existing_subnet = false

Configuration Parameter	Description	Sample Value
existing_subnet	This defines, if existing subn	external_existing_subnet = false

Configuration	Description	Sample Value
management_existing_subnet	This defines, if existing subn	management_existing_subnet = false

Configuration	Description	Sample Value
create_management_nat	This defines, if NAT needs to	create_management_nat = true

Configuration	Description	Sample Value
create_internal_nat	This defines, if NAT needs to	create_internal_nat = true

Configuration	Description	Sample Value
create_external_ig	This defines, if internet gate	create_external_ig = true

Configuration	Description	Sample Value
use_static_external_ip	This defines, if static private	use_static_external_ip = false

Configuration	Description	Sample Value
use_static_internal_ip	This defines, if static private	use_static_internal_ip = false

Configuration	Description	Sample Value
use_static_management_ip	This defines, if static private	use_static_management_ip = false

Configuration	Description	Sample Value
use_load_balancer	This defines, if load	use_load_balancer = false

		Description	**Sample Value**	
		balancer needs to be		

		Description	**Sample Value**	
		all interface.		

Configuration	Description	Sample Value
use_existing_lb_for_ext_nic	This defines, if existing load	use_existing_lb_for_ext_nic = false

Configuration	Description	Sample Value
use_existing_bs_for_ext_nic	This defines, if existing bback	use_existing_bs_for_ext_nic = false

Configuration	Description	Sample Value
use_existing_ls_for_ext_nic	This defines, if existing list	use_existing_ls_for_ext_nic = false

Configuration Option	Description	Sample Value
compartment_id	The OCID of the compartment wh	compartment_id = "ocid1.tenancy.oc1..aaaaaaaarod5yc3653ujwgvjbqui3s6r6ntfbs3d4uwxlkitku5flcbkrety"

Configuration Parameter	Description	Sample Value
2v0_display_name	Display name for the nZTA gate	vm_display_name = "skrn-new-2"

Configuration Parameter	Description	Sample Value
image_id	OCID of the image to use for t	image_id = "ocid1.image.oc1.ap-hyderabad-1.aaaaaaaapw7vnd4fqbp3heuk5kikfhlmhpcxjhcrl7tz7a3ln52gg7jr3h

Configuration Parameter	Description	Sample Value
2shape	VM shape for the nZTA gateway	shape = "VM.Standard.E4.Flex"

Configuration	Description	Sample Value
total_flex_OCPUs	Number of requested OCPU's for	total_flex_OCPUs = 3

Configuration	Description	Sample Value
total_flex_RAM	RAM size for the flex VM	total_flex_RAM = 16

Configuration	Description	Sample Value
availability_domain	Availability domain where the	availability_domain = "bsUY:AP-HYDERABAD-1-AD-1"

Configuration	Description	Sample Value
internal_vcn_cidr_block	Internal CIDR block for the	internal_vcn_cidr_block = "10.0.0.0/16"

Configuration	Description	Sample Value
internal_vcn_display_name	Display name for the internal VCN.	internal_vcn_display_name = "ZTA_internal_vnc_2"

Configuration Parameter	Description	Sample Value
internal_vcn_id	OCID of the existing vcn to use	internal_vcn_id = "ocid1.vcn.oc1.ap-hyderabad-1.amaaaaaatgkbhxyaia3zy75gt3cqxqotgdagfsw7vdy2sylscjvavm2

Configuration Parameter	Description	Sample Value
internal_subnet_id	OCID of the subnet to use for	internal_subnet_id = "ocid1.subnet.oc1.ap-hyderabad-1.aaaaaaaafscfkn7nzwsnv5xgjscybresi55fyxreqrqeu67umi

Configuration	Description	Sample Value
internal_subnet_cidr_block	CIDR block for the internal subnet	internal_subnet_cidr_block = "10.0.0.0/18"

Configuration	Description	Sample Value
internal_subnet_display_name	Display name for the internal subnet	internal_Subnet_2"

Configuration	Description	Sample Value
internal_rt_name	Display name of the routing table	internal_rt_name = "internal_rt_name"

Configuration	Description	Sample Value
internal_nat_name	Display name of the nat for the	internal_nat_name = "internal_nat_name"

Configuration	Description	Sample Value
internal_nic_display_name	Display name for the internal NIC	internal_nic_display_name = "internal"

Configuration	Description	Sample Value
internal_ip_address	Static IP address of the internet	internal_ip_address = "10.1.1.2"

	Configuration	Description	Sample Value
36	internal_is_public_ip_enabled	This defines, if we need to ss	internal_is_public_ip_enabled = false

Configuration	Description	Sample Value
internal_nsg_name	Display name of the network service	internal_nsg_name = "internal_nsg_name"

Configuration	Description	Sample Value
internal_nsg_rules		protocol = "6"//TCP destination ="0.0.0.0./0" source = "0.0.0.0./0" min_dstport = 6667 max_dstport = 6667 min_srcport = 0 max_srcport = 0 description = "internal ingess" direction = "INGRESS"

Configuration Variable	Description	Sample Value
external_vcn_cidr_block	External CIDR block for the external	External_vcn_cidr_block = "10.0.0.0/16"

Configuration	Description	Sample Value
external_vcn_display_name	Display name for the external VCN	external_vcn_display_name = "ZTA_external_vnc_2"

Configuration Parameter	Description	Sample Value
external_vcn_id	OCID of the existing vcn to use	external_vcn_id = "ocid1.vcn.oc1.ap-hyderabad-1.amaaaaaatgkbhxyaia3zy75gt3cqxqotgdagfsw7vdy2sylscjvavm

	Configuration	Description	Sample Value
42	external_subnet_id	OCID of the subnet to use for	external_subnet_id = "ocid1.subnet.oc1.ap-hyderabad-1.aaaaaaaafscfkn7nzwsnv5xgjscybresi55fyxreqrqeu67umi

Configuration Name	Description	Sample Value
external_subnet_cidr_block	CIDR block for the external subnet	external_subnet_cidr_block = "10.0.0.0/18"

Configuration	Description	Sample Value
external_subnet_display_name	Display name for the external subnet	external_Subnet_2"

Configuration	Description	Sample Value
external_rt_name	Display name of the routing table	external_rt_name = "external_rt_name"

Configuration Option	Description	Sample Value
external_ig_name	Display name of the external ig	external_ig_name = "external_ig_name"

Configuration	Description	Sample Value
external_nic_display_name	Display name for the external NIC	external_nic_display_name = "external"

Configuration	Description	Sample Value
external_ip_address	Static IP address of the external	external_ip_address = "10.1.1.2"

	Configuration	Description	Sample Value
49	external_is_public_ip_enabled	This defines, if we need to ss	external_is_public_ip_enabled = true

Configuration	Description	Sample Value
5e0xternal_nsg_name	Display name of the network se	external_nsg_name = "external_nsg_name"

Configuration	Description	Sample Value
5e1xternal_nsg_rules		protocol = "6"//TCP destination ="0.0.0.0./0" source = "0.0.0.0./0" min_dstport = 6667 max_dstport = 6667 min_srcport = 0 max_srcport = 0 description = "external ingess" direction = "INGRESS"

	Configuration	Description	Sample Value
52	management_vcn_cidr_block	management CIDR block for the	management_vcn_cidr_block = "10.0.0.0/16"

Configuration	Description	Sample Value
management_vcn_display_name	Display name for the management VCN	management_vcn_display_name = "ZTA_management_vnc_2"

Configuration	Description	Sample Value
management_vcn_id	OCI Domain the existing vcn to use	management_vcn_id = "ocid1.vcn.oc1.ap-hyderabad-1.amaaaaaatgkbhxyaia3zy75gt3cqxqotgdagfsw7vdy2sylscj

Configuration	Description	Sample Value
management_subnet_id	OCI Domain the subnet to use for	management_subnet_id = "ocid1.subnet.oc1.ap-hyderabad-1.aaaaaaaafscfkn7nzwsnv5xgjscybresi55fyxreqrqeu6

Configuration Variable	Description	Sample Value
management_subnet_cidr_block	CIDR block for the management subnet	management_subnet_cidr_block = "10.0.0.0/18"

Configuration Name	Description	Sample Value
management_subnet_display_name	Display name for the management subnet	management_subnet_display_name = "management_Subnet_2"

Configuration	Description	Sample Value
management_rt_name	Displays the name of the routing table	management_rt_name = "management_rt_name"

Configuration	Description	Sample Value
management_nat_name	Displays the name of the nat for the	management_nat_name = "management_nat_name"

Configuration	Description	Sample Value
management_nic_display_name	Display name for the management interface	management_nic_display_name = "management"

Configuration	Description	Sample Value
management_ip_address	Static IP address of the management	management_ip_address = "10.1.1.2"

		Configuration	Description	Sample Value
	62	management_is_public_ip_enable	This defines, if we need to ss	management_is_public_ip_enabled = false

Configuration	Description	Sample Value
management_nsg_name	Display name of the network security group.	management_nsg_name = "management_nsg_name"

Configuration	Description	Sample Value
6m4a_management_nsg_rules		protocol = "6"//TCP destination ="0.0.0.0./0" source = "0.0.0.0./0" min_dstport = 6667 max_dstport = 6667 min_srcport = 0 max_srcport = 0 description = "management ingess" direction = "INGRESS"

	Configuration	Description	Sample Value
65	init_config	Init config that nZTA will pic	init_config = "PHB1bHNlLWNvbmZpZz48cHJpbWFyeS1kbnM+OC44LjguODwvcHJpbWFyeS1kbnM+PHNlY29uZGFyeS1kbnM pbj48Y2VydC1jb21tb24tbmFtZT5za3JuLW9jaS5nLmUyZTluZS5qdW5pcGVyLnB6dC5kZXYucGVyZnNlYy5jb208L2 W5zZS1hZ3JlZW1lbnQ+PGNvbmZpZy1kb3dubG9hZC11cmw+J2h0dHBzOi8vZTJlMi5qdW5pcGVyLnB6dC5kZXY UzNDMzYi9vcmNoZXN0cmF0aW9uL2luaXRpYWwtY29uZmlnP3Q9Z0FBQUFBQmstQUFmYXIwYjgtNXowMXo4aV QWp2c2JEWGUzcnA1X2VmU3Joc1lSWTM1SU96WmE3dlZsaDRXalNQ3ZXa3JFVkVVek5mOTdqeDI1T1A1VktxQ pYWDdOa2Y5STZlWUctUWI0aHRENzNOZnZyYjhzNmRYZGo0WUlIaXItdXA0YnJ5d1I0dUdiVThuZ20zR3ZWanFPa ZTUFmY1E4YTM5MmRFUlQ1OVhWM3p4QkRvaklQQ1l1RlNtbExCd2NKckRjZC1oTVMwSmpxYUE1NHhLMDNsM UF4TT0nPC9jb25maWctZG93bmxvYWQtdXJsPjxhcHBsaWFuY2UtaWQ+MGNhM2U2ZTBmZjEyNDUxN2FhNjNjO XlucHp0LmRldi5wZXJmc2VjLmNvbTwvY29udHJvbGxlci1ob3N0bmFtZT48Y29udHJvbGxlci1lbnJvbGxlZC1ob3N0b bGVkLWhvc3RuYW1lPjxkbnMtc2VhcmNoLWRvbWFpbj5wc2VjdXJlLm5ldDwvZG5zLXNlYXJjaC1kb21haW4+PGN PjwvcHVsc2UtY29uZmlnPg=="

Configuration	Description	Sample Value
lb_display_name	Display name of the load balancer	lb_display_name = "external_lb_name"

Configuration	Description	Sample Value
listener_display_name	Display name of the listener	listener_display_name = ["external_ls_name_443" , "external_ls_name_80"]

Configuration	Description	Sample Value
bs_display_name	Display name of the backend server.	bs_display_name = ["external_bs_name_443","external_bs_name_80"]

Configuration	Description	Sample Value
external_lb_id	OCI id of the external load bal	external_lb_id = "ocid1.networkloadbalancer.oc1.ap-hyderabad-1.amaaaaaatgkbhxyanzrjec4irpjuumytyuk5qugc

Configuration	Description	Sample Value
ports	Load balancer ports for the li	ports = [443,80]

Configuration	Description	Sample Value
7b1s_policy	Backend set policy that should	bs_policy = "FIVE_TUPLE"

Configuration	Description	Sample Value
protocol	TCP is only protocol supported	protocol = "TCP"

Creating a VM Instance of the Uploaded OCI Image Using any Other Methods

Deploy a VM with nZTA Gateway image uploaded to the OCI, with following requisites:

1. Configure 2 nic's in the order Internal, External, if the Management port is not required for the nZTA Gateway deployment.
2. Configure 3 nic's in the order Internal, External and Management. If management port is configured for the nZTA Gateway deployment.
3. Ensure custom metadata "pulse-config" is configured for the VM. The value of pulse-config metadata needs to be taken from the init file downloaded from the nZTA Controller interface as explained in admin guide section "Downloading Metadata for Oracle Cloud Platform" on page 335.
4. Ensure internal NIC can access the controller present in the azure cloud and application resources.
5. Ensure external NIC public ip is reachable by the ISAC clients over the port 443.
6. Ensure management NIC can access the controller present in the azure cloud , if management port is enabled for controller communication.
7. Ensure load balancer ip is reachable by the ISAC clients over port 443, if load balancer is used for gateway deployment.
8. Recommended Firewall rules:
    - Internal - INGRESS (TCP: 6667), EGRESS (TCP: ANY, UDP: ANY)
    - External - INGRESS (TCP: 443)
    - Management - INGRESS (TCP: 6667), EGRESS (TCP: ANY, UDP: ANY)

Upgrading Gateways

Ivanti periodically creates and releases new Gateway software versions to address updates and issues, and to improve performance. As new version packages become available, you can trigger an upgrade for your Gateways through the *Controller* to take advantage of the updates available in the new version. Gateway updates can be applied manually, or, in the case of ungrouped Gateways, applied at a scheduled time.

The **Installation Packages** page displays a list of the available Gateway installation packages and controls the update schedule. To trigger a manual upgrade for a specific Gateway or Gateway Group, use the **Gateways Overview** page (see Checking a Current Gateway Version and Applying an Individual Update).

The upgrade process requires a Gateway instance to become unavailable for a short time while the upgrade package is applied. *Ivanti* recommends performing Gateway upgrades ONLY at a time of least impact to your services.
Gateways that are part of a group (for high availability) can be safely upgraded provided the remaining group members remain connected and available.

Upgrading the ZTA Gateway to the versions 22.7R1 or 22.6R1.8 or 22.5R1.7 is recommended for the ZTA Controller version 22.7R1. Please be advised that there is no download package option available for versions 22.6R1.8 and 22.5R1.7. As an alternative, ZTA Gateways with the versions 22.5R1 or 22.6R1.2 can be downloaded and subsequently upgraded to the versions 22.5R1.7, 22.6R1.8 and 22.7R1.

To view the list of available Gateway installation packages and to configure a update schedule:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Administration** icon, then select **Upgrade > Installation Packages**.

   The *Installation Packages* page appears.

3. Click the **Gateways** tab.

   The list of available Gateway installation packages is displayed:

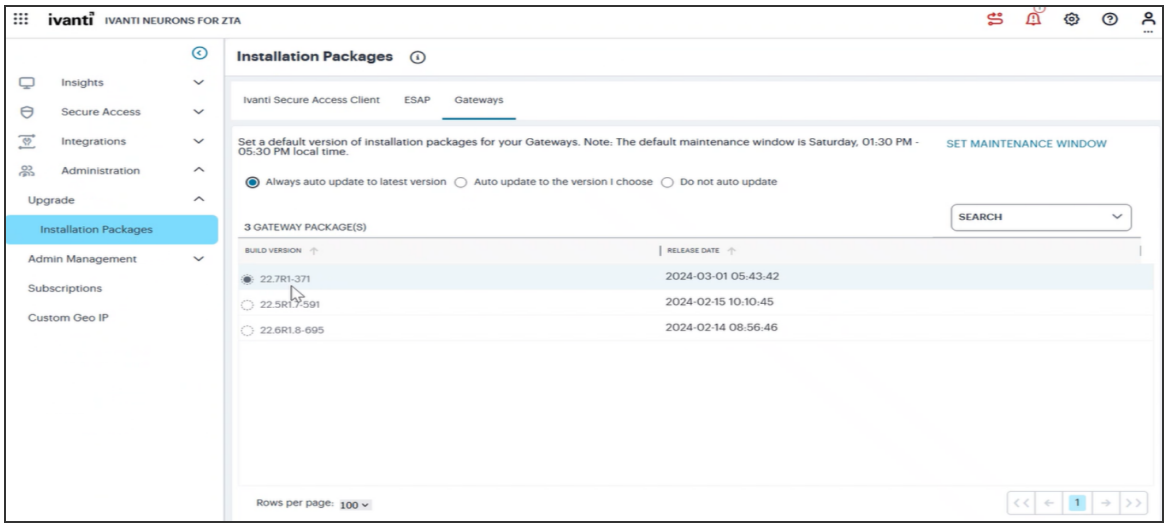

   View a list of available Gateway installation packages and set a schedule for upgrades

   Use the following controls to set the Gateway upgrade policy:

   - **Always auto update to the latest version**: Update your Gateways to the latest available version according to the defined maintenance window.
   - **I will choose the version**: Select a specific package from the list of available packages to be applied to your Gateways according to the defined maintenance window.
   - **Do not set a default version**: Do not automatically schedule an update, and leave your Gateways running the current version.

4. To set the default maintenance window for Gateway updates, click **SET MAINTENANCE WINDOW**.

   The Maintenance Window dialog appears:

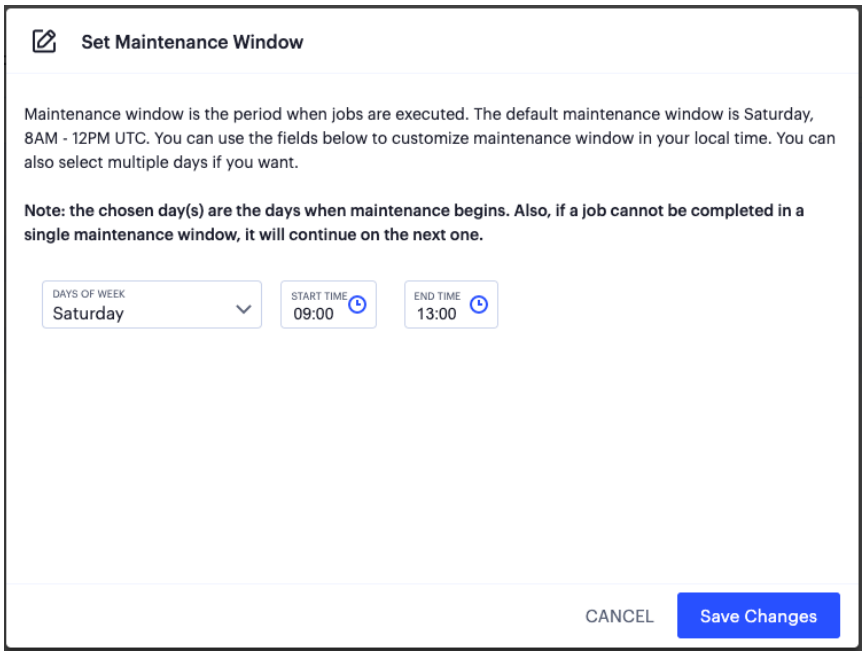

   Set the default maintenance window for Gateways

   Use the settings provided to configure your maintenance window, during which your non-grouped Gateways are updated to the selected package version. *Ivanti* recommends setting this at a time of least convenience to your services.

> If a scheduled update does not complete within the maintenance window, it continues at the next available maintenance window. However, if a scheduled update fails, try updating the Gateway manually (for more information, see Checking a Current Gateway Version and Applying an Individual Update).

Checking a Current Gateway Version and Applying an Individual Update

To check the current version for a Gateway, and to apply an update:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways > Gateway List**.

   The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*.
3. Select the required Gateway from the list.

   The *Gateways Overview* page appears. The summary at the top of the page displays details pertaining to this Gateway, including the current version:

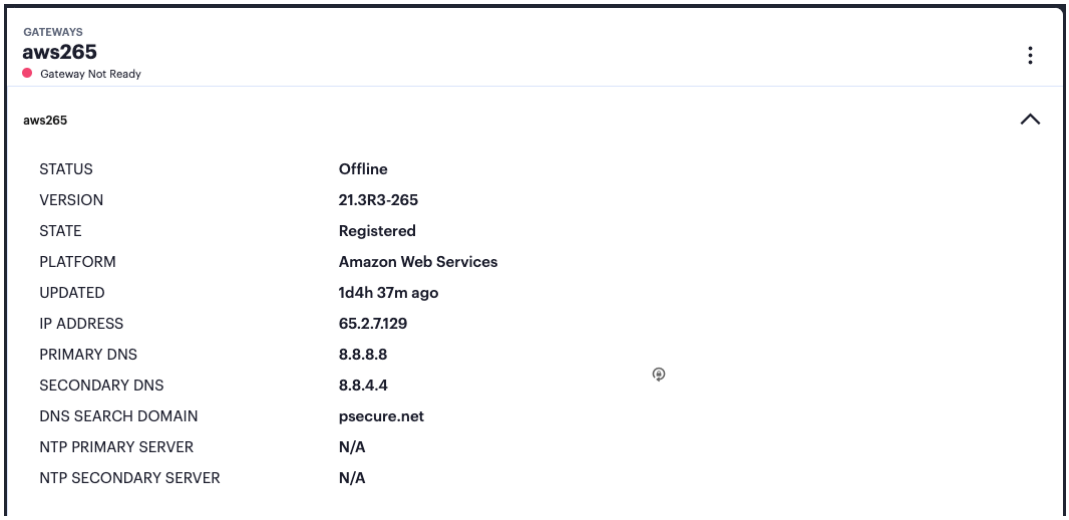

Viewing the current Gateway version

4. Click the context menu icon at the top-right to access the Edit options applicable to the selected Gateway:

   If an update is available for this Gateway, a corresponding link is displayed in the drop-down menu:

   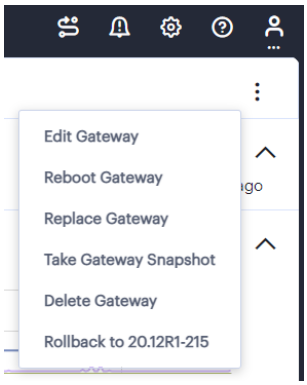

   Viewing Gateway version upgrade options

   > In some cases, there might be more than one version available. Select the version you want, or contact your support representative for details.

5. To start the upgrade, click the **Upgrade to <version>** link.

An upgrade task is added in the **Tasks** page for this Gateway, with a status of "Pending". As the task becomes due, the upgrade process starts and the task status changes to "In Progress". The status description describes the current phase of the upgrade, *Downloading*, *Installing*, or *Rebooting*, together with the percentage complete.

During the first two phases, your Gateway remains operating on the current version and continues to serve traffic. Then, in the reboot phase, the Gateway becomes unavailable for a short time.

After the reboot is complete, the Gateway automatically re-establishes connection to the *Controller*. If the procedure is successful, the upgrade task is marked with a status of "Success" and the new software version is displayed in the Gateway summary in the **Gateways Overview** page.

The time taken for an upgrade to complete is based on factors such as the number of previous pending tasks and network connection speed. Check on progress using the status description or through the host platform management console.

The *Controller* UI displays the upgrade phase progress indicator (downloading, installing or rebooting) only for Gateway upgrades from base version 21.2 upwards. For base Gateway versions earlier than 21.2, an indication of the current progress is available only from the host platform management console.

To troubleshoot issues related to upgrading your Gateway, view the Event Log (see Checking the Logs). Furthermore, if a Gateway fails to initiate a reboot at the end of the upgrade process, but is still accessible, check the log file on the Gateway instance using the following steps:

1. Use `ssh` to access the Gateway instance command line.
2. Locate and open the file `/tmp/doUpgrade.log`.

Rolling Back a Gateway to a Previous Version

Your *ZTA Gateway*s can be rolled back to a previously-installed version through the Tenant Admin Portal. You might want to return to an earlier version if, for example, you encounter an unforeseen issue with a newly-upgrading Gateway instance, or for testing purposes.

You can roll back a Gateway version only where that Gateway instance has been previously upgraded through the Tenant Admin Portal, and only to the previously-installed version.

To roll back a Gateway to an earlier version:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Gateways > Gateway List**.

   The *Gateways List* page appears, showing the full list of Gateway Groups and standalone Gateways currently configured on the *Controller*.
3. Select the required Gateway from the list.

   The *Gateways Overview* page appears. The summary at the top of the page displays details pertaining to this Gateway, including the current version.

4. Click the context menu icon at the top-right to access the Edit options applicable to the selected Gateway:

   If a rollback function is available for this Gateway, a corresponding link is displayed in the drop-down menu:

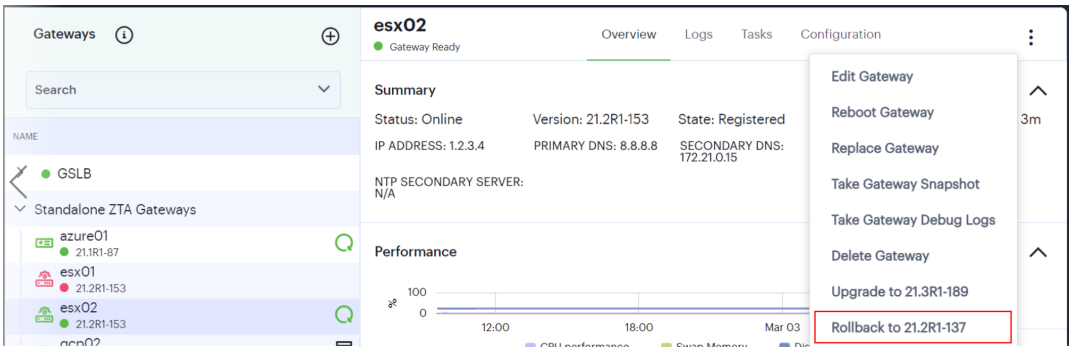

   Viewing the Gateway version rollback option

5. Select the rollback link.

   A rollback task is added in the **Tasks** page for this Gateway, with a status of "Pending". As the task becomes due, the rollback process starts and the task status changes to "In Progress".

   After the reboot is complete, the Gateway automatically re-establishes connection to the *Controller*. If the procedure is successful, the upgrade task is marked with a status of "Success" and the new software version is displayed in the Gateway summary in the **Gateways Overview** page.

   As the process starts, your Gateway remains operating on the current version and continues to serve traffic. Then, after the earlier version is reinstated, the Gateway reboots and becomes unavailable for a short time.

Configuring a Default Gateway for Application Discovery

nZTA directs requests from each application towards the Gateway defined in the secure access policy for the application.

For requests for applications not referenced by a secure access policy, you can define a *default Gateway* on the *Controller*. This enables packet analysis to be conducted on requests passing through the Gateway to assess the validity of the requests.

To view usage metrics for application and resource requests handled by the default Gateway, see Reviewing Application Usage.

Two default Gateway scenarios are supported:

- Any single Gateway at v21.1 (or later) can be assigned to act as the default Gateway. This Gateway is used exclusively as the default Gateway.
- Alternatively, any Gateway Group whose Gateways are all at v21.1 (or later) can be assigned to act as the default Gateway. In this scenario, all Gateways in the group are used exclusively as the default Gateway. The Gateway Group is typically fronted by a load balancer to enable the required distribution of requests across the Gateways in the group.

To configure a default *nZTA Gateway*, you must edit and update the built-in *Application discovery* secure access policy to reference a single Gateway or Gateway Group.

The default Gateway (or Gateway Group) then handles all requests from applications on enrolled devices that are not referenced by any other secure access policy.

Ivanti Secure Access Client version 21.1 is required to work with application discovery and a default Gateway.

Ivanti Secure Access Client Linux variants do not currently support the use of a default gateway.

To assign a Gateway (or Gateway Group) as the default Gateway:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Secure Access Policies**.

   The *Secure Access Policies* page appears. This lists all current secure access policies.
3. Locate the *Application discovery* secure access policy and examine its **Enabled** property:
   - If the policy is enabled, select the check box for the policy and then click **Disable**.

4. Select the check box for the *Application discovery* secure access policy and click **Edit**.

   The **Edit Secure Access Policy** dialog appears.

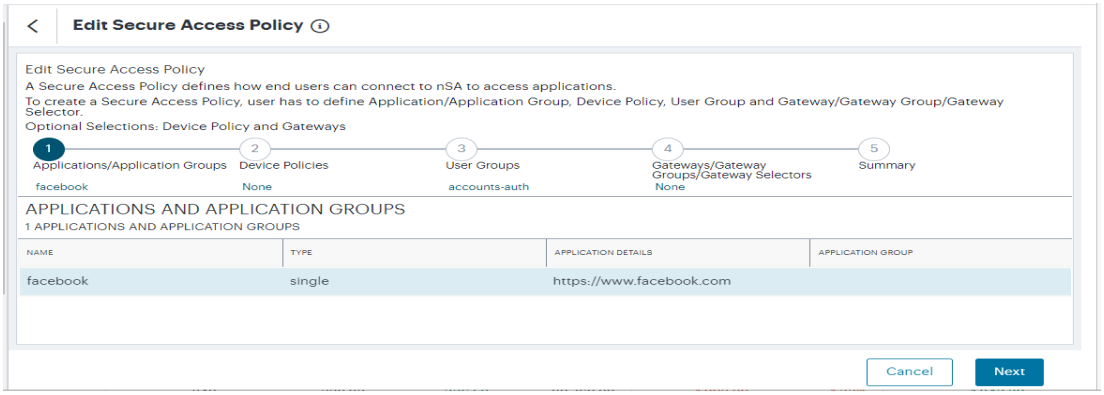

   Editing a Secure Access Policy

   The **Application Type** and **Application** cannot be changed.

5. Click **Device Policy** and select the required device policy.
6. Click **User Group** for the required user group. All users identified in this group will be able to access the default Gateway.
7. (Optional) For a default Gateway that is a single Gateway:
   - Ensure **Gateway Type** is set to *Single*.
   - Click **Select a Gateway** and select the required default Gateway.

   This Gateway must be at (v21.1 or later), and cannot already be referenced by another secure access policy.

   - Click **Save**. The **Secure Access Policies** page updates to reflect the selection.

8. (Optional) For a default Gateway that is a Gateway Group:
    - Ensure **Gateway Type** is set to *Group*.
    - Click **Select a Gateway Group** and select the required default Gateway Group.

    > All Gateways in this Gateway Group must be at (v21.1 or later), and cannot already be referenced by any other secure access policies.

    - Click **Save**. The **Secure Access Policies** page updates to reflect the selection.
9. Select the check box for the *Application discovery* secure access policy and click **Enable**.

    The **Secure Access Policies** page updates to reflect the selection. The policy is pushed to the selected Gateway or Gateway Group.

The default gateway configuration is complete, and the default Gateway can be used by any enrolled macOS/Windows desktop device to support application discovery after it accepts the new policy.

To enroll and use *Ivanti Secure Access Client* on a device, see Enrolling Mobile/Desktop Clients.

Configuring nZTA Gateway Connection Control for Trusted Networks

nZTA Gateway can sometimes be bypassed so that users can connect directly to specific applications. For example, you might want users to bypass nZTA for a specific application if they are connected directly to your trusted corporate network. nZTA Gateway tunnel creation will be bypassed on the endpoint since resource access will go through the physical interface.

In the Gateway Selector, a network tag must be created for Known Networks (On-Prem) and this network tag must be marked to bypass the default gateway in the Gateway Selector so that all traffic from known network will be bypassed.

> Device Posture evaluation is not done when nZTA Gateway is bypassed.

> No analytics data will be displayed on any dashboards when nZTA Gateway is bypassed.

To configure bypass settings for nZTA:

1. Create Network Tags for Known Network. For example, On-Prem.

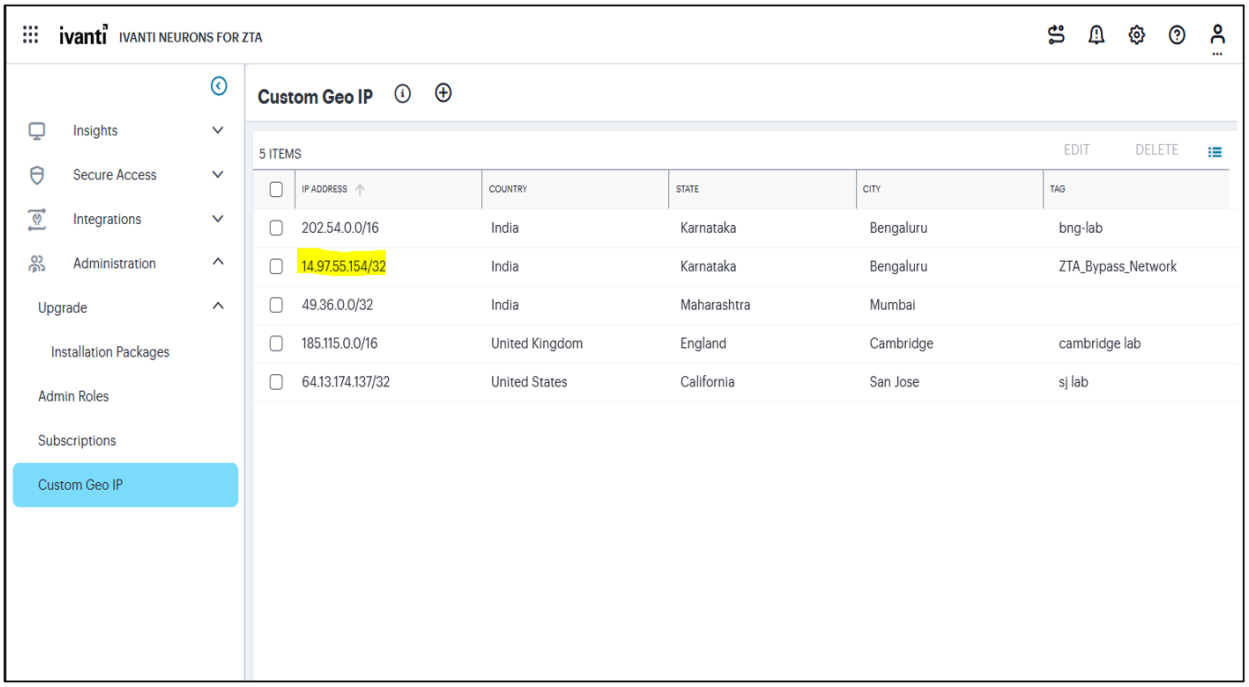

2. Click **Manage Gateways**, under Gateway Selectors, enter the Gateway information, Add Rules for Recognized Networks.

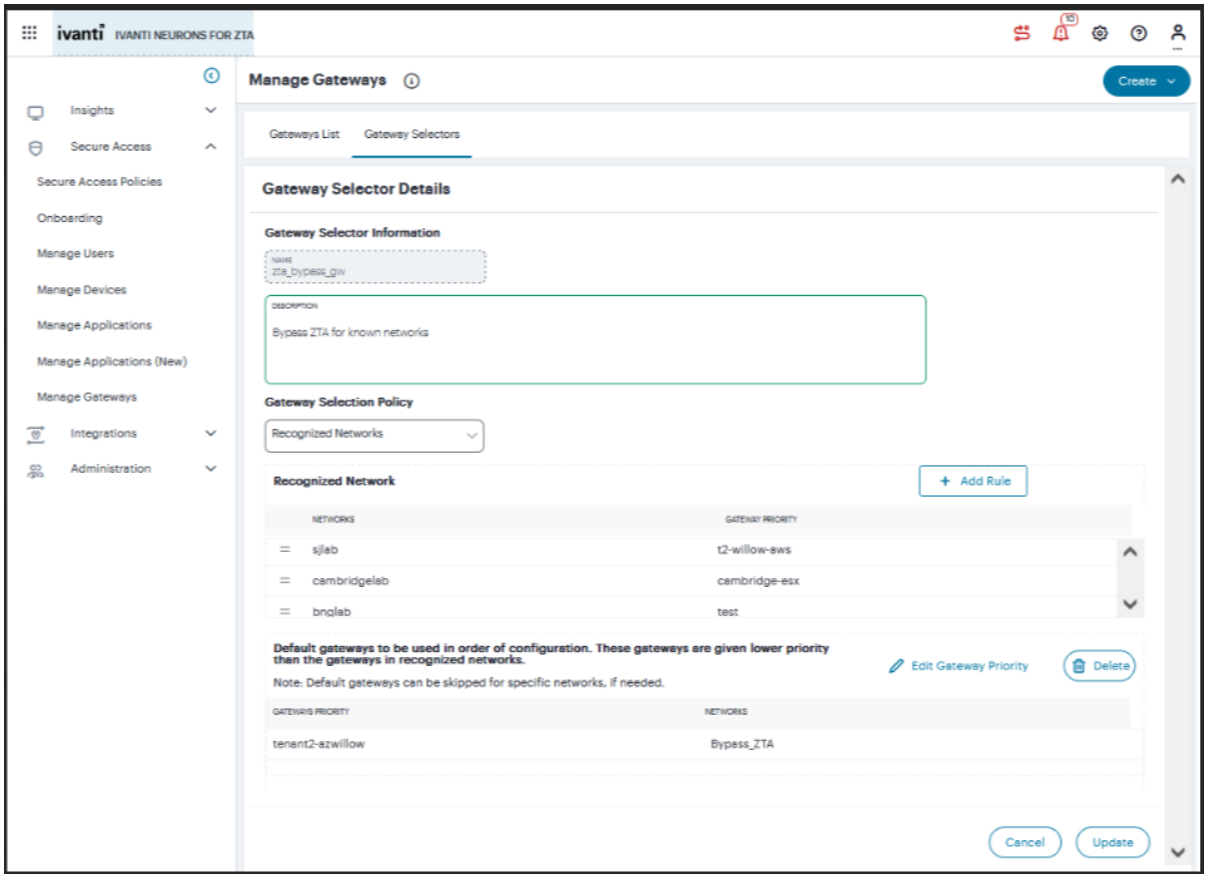

3. Specify the default gateways to be skipped if needed and click **Update**. The traffic will be bypassed through nZTA Gateway and it will be passed through physical interface.

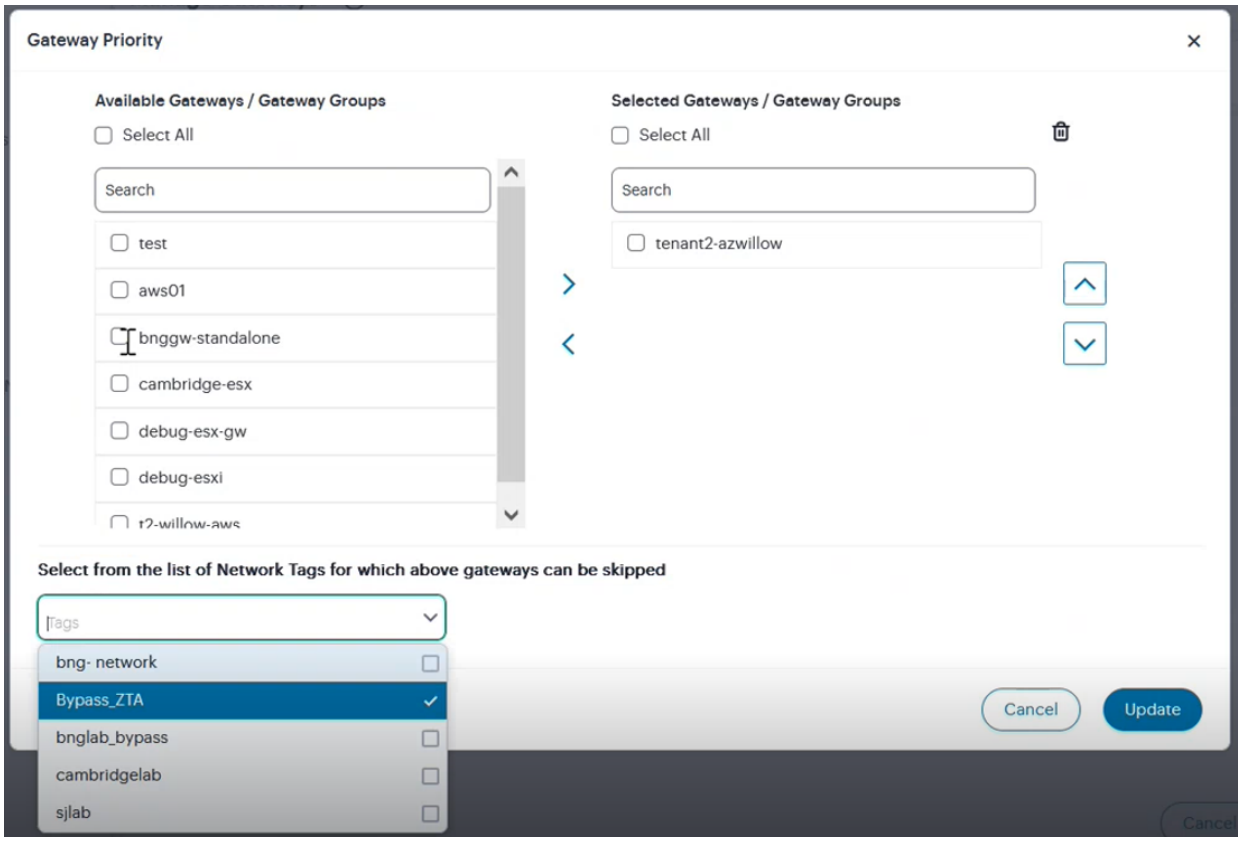

Creating Device Policies and Device Policy Rules

Introduction

Device Policies define how desktop and mobile devices access cloud and on-premise applications in your *Ivanti Neurons for Zero Trust Access* (*nZTA*) deployment.

> Device policies act as one of the four dimensions of a Secure Access Policy, see Creating/Editing Secure Access Policies.

You create a device policy and then create / associate the device rules to form a complete **Device Policy**, suitable for adding to a **Secure Access Policy**. Device policies encompass a set of rules that define the minimum standard a device must meet to be considered compliant with the applications and services served by your Secure Access Policies.

To learn more about Device Policies and Rules, see "Viewing Device Policies and Rules" below.

Viewing Device Policies and Rules

nZTA provides a number of built-in default device policies, each containing a set of appropriate built-in device rules. You cannot modify / delete these built-in default device policies. These policies and rules are suitable for general use. In addition, *nZTA* allows the definition of custom policies and rules to fit an organization's specific requirements.

To view the list of all default and custom device policies or rules defined on the *Controller*:

1. Log into the *Controller* as a Tenant Admin.
2. From the *nZTA* menu, select the **Secure Access** icon, then select **Manage Devices > Device Policies**.

   The *Device Policies* page appears. This page lists all current device policies.

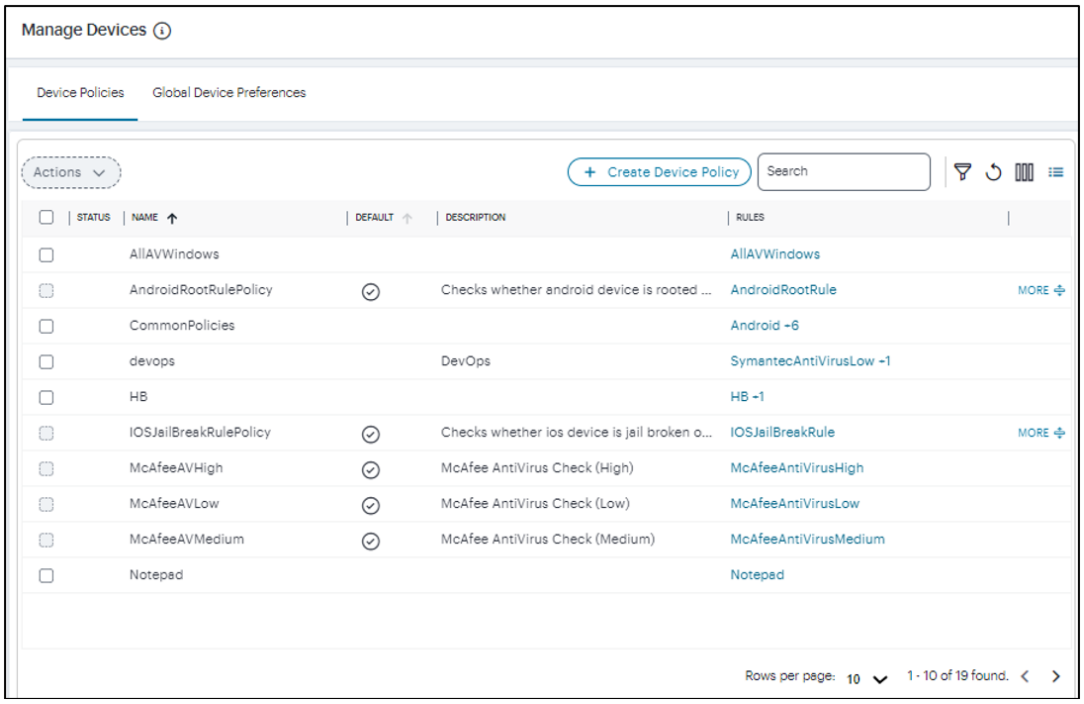

   The Device Policies page

   Built-in default policies are indicated by a tick in the **Default** column. Custom policies are not ticked.

   On this page, you can:

   - Add a new custom device policy, see Creating Device Policies.
   - Edit / Delete custom device policy, see "Editing / Deleting Custom Device Policy" on page 447.
   - Edit / Delete custom device policy rule, see "Editing / Deleting a Custom Device Policy Rule" on page 455.

You can also:

- Sort the list by a selected column in ascending or descending order.
- Switch between normal and denser data views.

Creating Device Policies

You can create **Device policies** and then create / associate one or more **Device Rules** as required.

To create a device policy:

1. Log into the *Controller* as a Tenant Admin.
2. From the *nZTA* menu, select the **Secure Access**, then select **Manage Devices > Device Policies**.

   The *Device Policies* page appears. This page lists all current device policies.
3. Click **Create Device Policy**.

   A form appears to enable you to create the device policy.

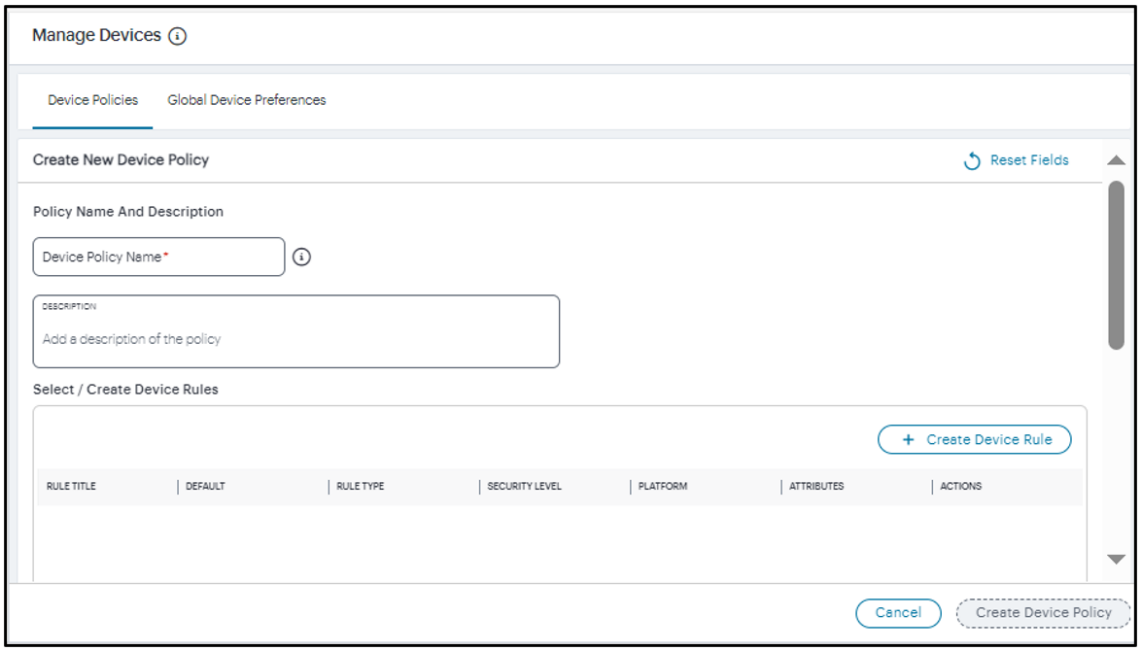

Create a new Device Policy

> At any point during this process, you can reset the form data by selecting **Reset Fields**.

4. Enter a **Name** for the device policy.
5. Add a **Description** for the device policy.

6. Select each of the listed **Rules** that are required for the device policy, or select **Create Device Rule** to use the in-line rule creation form. To learn more about this process, see Creating Device Policy Rules .

7. (Optional) In the *Rule Requirement* section: Specify for each end-user device **Platform** how you want to enforce your policy rules by choosing one of the following **Rule Requirement** options:

    - **All of the above rules**: The end-user device must comply with all rules defined in the policy.
    - **Any of the above rules**: The end-user device must comply with at least one of the defined rules in the policy.

- **Custom**: The end-user device must comply with the conditions specified in a custom expression. Use the **Custom Expression** field to define an expression for the rules defined in this policy and how they should be evaluated. You can use the Boolean operators AND, OR and NOT, and also use parentheses to group or nest conditions.

  The following is a list of sample custom expressions:

  - *customExpr*
  - *(customExpr)*
  - *NOT customExpr*
  - *customExpr OR customExpr*
  - *customExpr AND customExpr*

  As an example, where a policy has associated with it the rules "Rule1", "Rule2", and "Rule3", the following expression is valid: *Rule1 AND (NOT Rule2 OR (NOT Rule3))*

  When using custom expressions, consider the following points:

  - Using NOT: When using "*NOT expr*", the negated expression evaluates to true if the outcome of *expr* is false and evaluates to false if the outcome of *expr* is true.
  - AND, OR, NOT precedence: These operators are evaluated from highest to lowest precedence in this order: NOT (from right), AND (from left), OR (from left).
  - A combination of any device rule is allowed in an expression, except location, time of day, and network rules. For example, the following expressions are not allowed:
    - Windows_Process AND Locationrule
    - Windows_Process AND Networkrule
    - Windows_Process AND Time-of-Day_Rule

After you have set a platform and rule requirement, select **Apply** to add the entry. Then, repeat this procedure if you want to add any rule requirements for other device platforms.

If you intend to add multiple rules of varying types to a device policy, be aware that individual rules might not by themselves guarantee allowed or denied access to an application depending on the outcome of other evaluated rules in a device policy, and the rule requirements settings configured here.

8. (Optional) In the *Remediation* section: To provide custom remediation instructions for the policy, tick **Enable Custom Instruction** and enter your remediation text into **Custom Instruction**. This option also requires selection of a target **Platform**.

   These instructions are presented through *Ivanti Secure Access Client* when a device compliance check fails based on this policy.

   - This feature is applicable to Windows, Mac, and Linux device policies only.
   - Also note that custom instructions are restricted to a 500 byte limit and can contain only plain text or an HTML document with HREF links.

9. Select **Create Device Policy**.

   The new device policy appears in the list of **Device Policies**.

10. Repeat steps 3-7 to create all required device policies.

After you have created all required device policies, you can move to the next stage of *nZTA* configuration, which is Creating/Editing Secure Access Policies.

Editing / Deleting Custom Device Policy

Built-in default device policies are indicated with a tick mark in the Devices column, and they cannot be edited or deleted.

To edit a custom device policy:

1. In the Device Policies page, select the check box next to the custom device policy that you want to edit.

2. Select **Actions > Edit**.

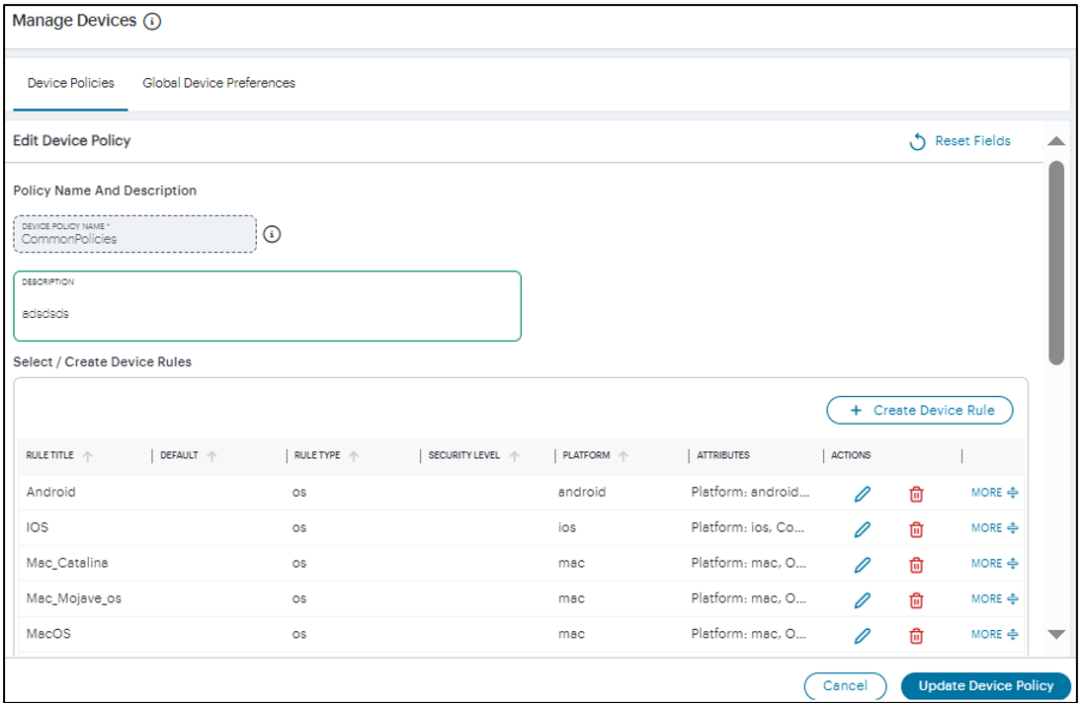

3. Make the necessary changes, such as change the policy description, or create / edit / delete rules. You are not allowed to change the policy name.

4. Click **Update Device Policy**.

To delete custom device policy:

1. In the Device Policies page, select one or more check boxes next to the custom device policies that you want to delete.

2. Select **Actions > Delete**.

3. In the delete confirmation window, click **Yes, Delete**.

Configuring Default Device Policy for Users

As part of configuring an application, we can associate a device policy, which may have one or more same or different type of device rules configured. When a user tries to login, AAA evaluate these policies, log failures and allows sign in. When a user tries to access applications, device policies are evaluated and enforced. If a device policy evaluation fails, application access is denied.

With Default device policy for users, Admins can configure policies that get enforced even before device authentication, that is during the user enrollment or user authentication.

- For default enrollment policy, User Group will always be added.
- For a new multi-sign-in policy of type enroll, always add User Group first with the new enroll policy.
- Visibility/analytics in the form of charts are not available, but logs are available in Insight >logs.
- Risksense policy when enforced on enrollment sign-in policy is not supported with web/browser based enrollment, but is supported when Ivanti Secure Access Client is already installed.
- Time of day rule type is not supported for default device policy.
- Time of day and OS check rules are not supported on the enrollment sign in url when trying to enroll from iOS endpoint.

You can use the existing default polices or can create new policy and use the default device policy.

To configure default Device Policy for users:

1. Log into the Controller as a Tenant Admin.
2. From the *nZTA* menu, select **Secure Access** and then select **Manage Users > User Policies**.
3. Click **Create User Policy**.

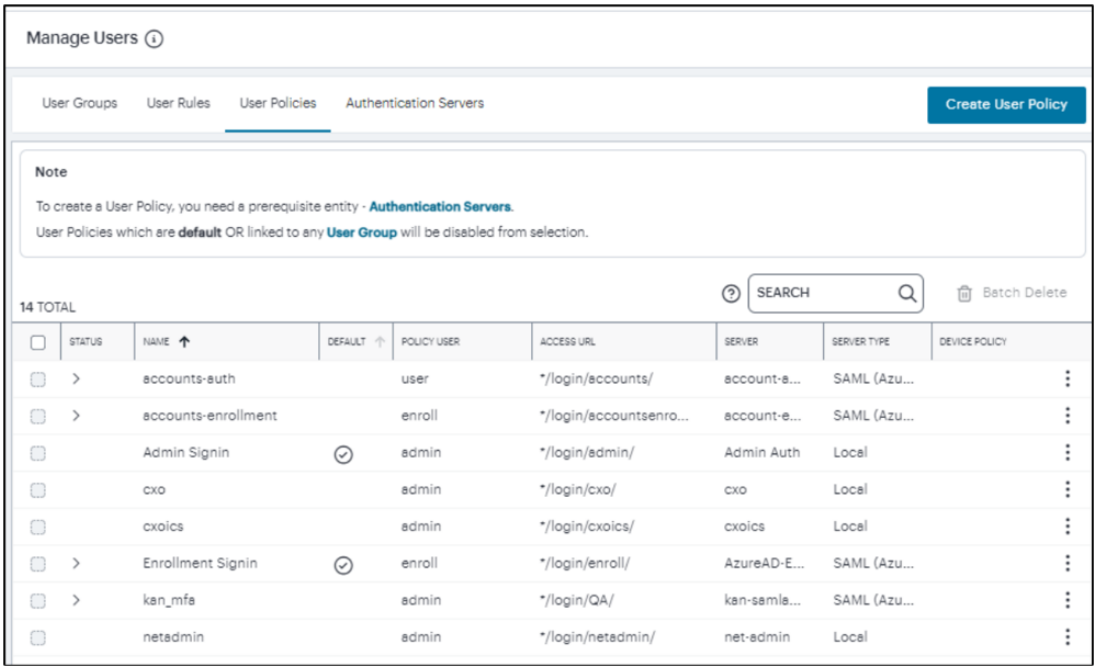

Manage User Policy

4. Enter the **Policy Name**, **Login URL** using the format */login/<path>.
5. Select the **User Type: Enrollment Users/ Users/Administrators**.
6. Select the **Device Policy** from the drop-down menu. For example, Deny_Location.

There are a few exceptions while creating User Policies when user Type is Administrator. The following device policies are not applicable to Administrator user.
- Any device policy having Risk Sense rule.
- Any device policy having Time of Day rule.
- Any device policy having combination of Location and Network rules.

Continuous Device Posture Assessment (CARTA) is not supported for Admin Access. However, the Device Posture Assessment will occur during the administrator login process, if configured.

7. Select the **Auth Server**.
8. Click **Create User Policy**.

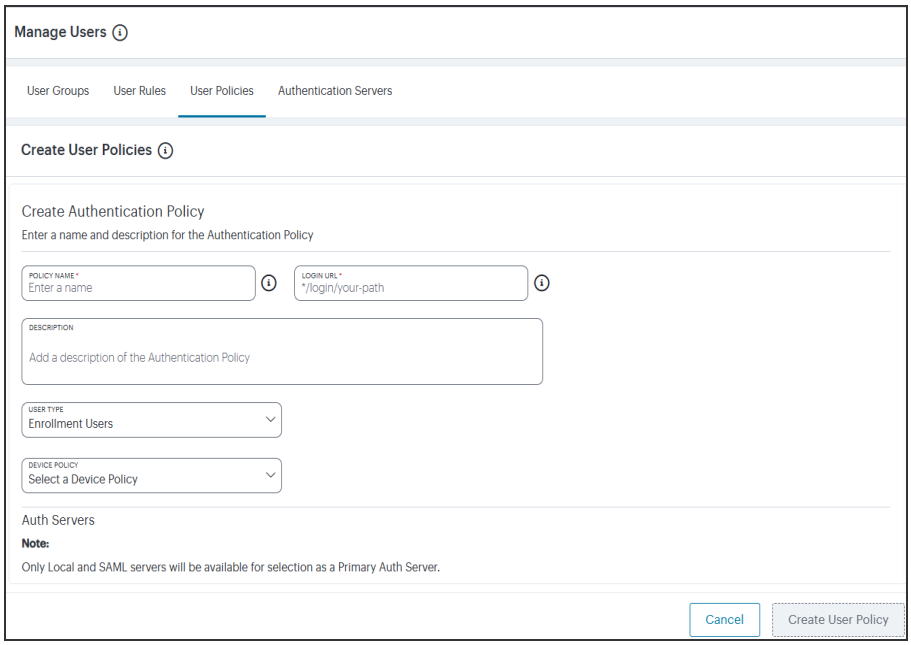

Create User Policy

9. Click **Create User Policy**.

10. Users can also edit the existing Default policy to include the Device policy during the enrollment sign-in/user authentication.

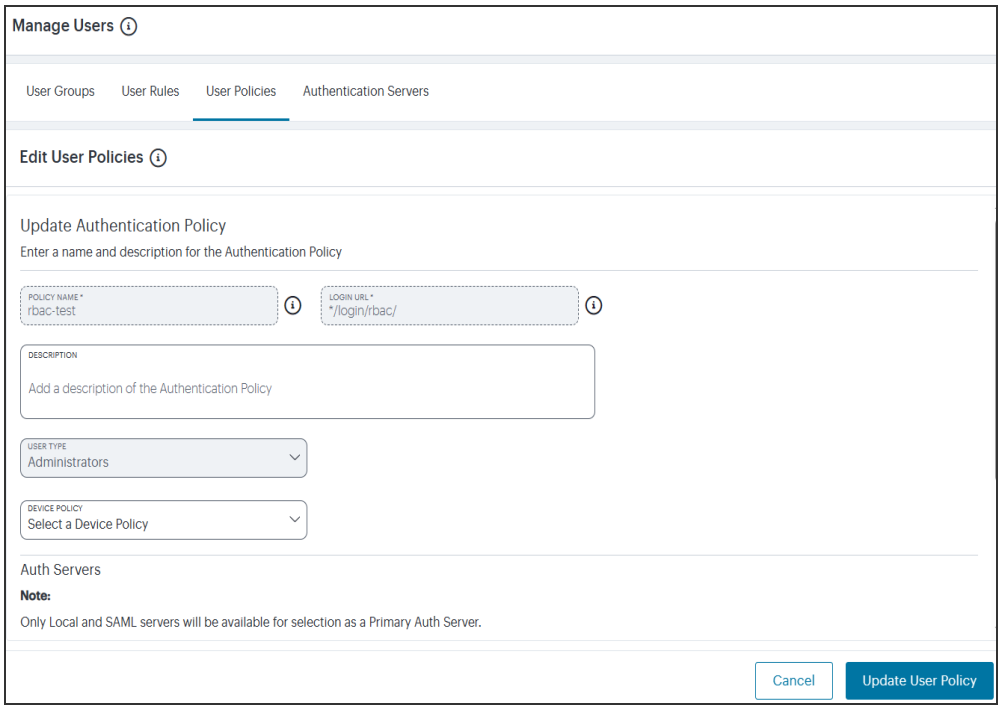

Edit User Policy

Creating Device Policy Rules

Before you begin, decide what kind of rule you want to create. For each rule type, make sure you have the supporting parameters. For example, if you are creating a *Network* rule, make sure you know the IP address and netmask range you want to apply.

To create a device rule:

1. In the Create Device Policy page, click **Create Device Rule**.

   The **Create Device Rule** form appears.

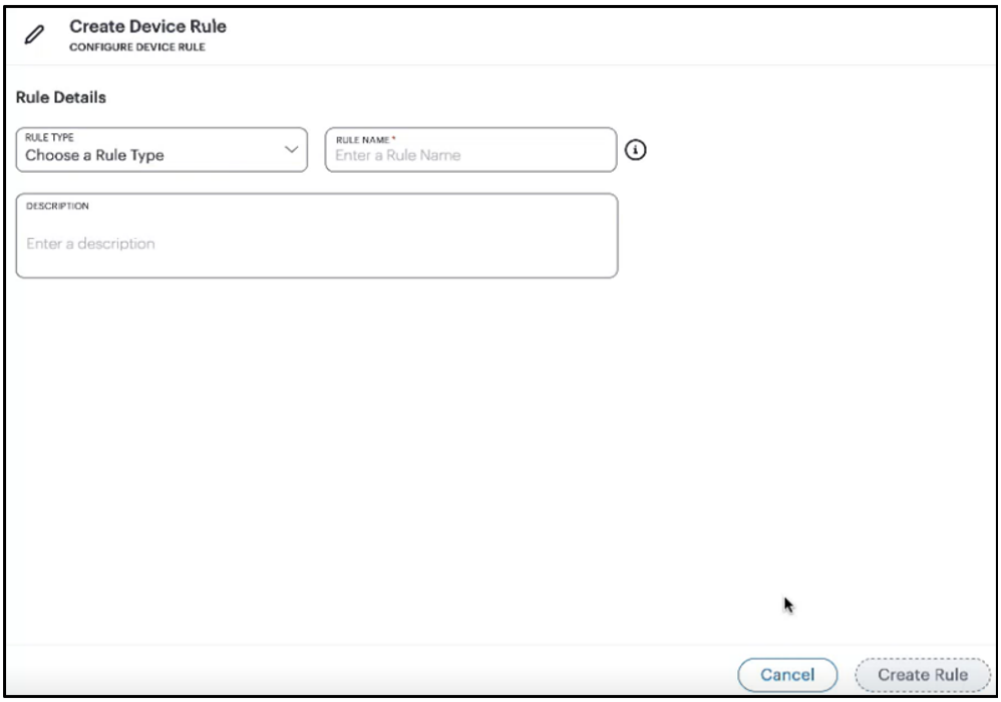

Create Device Rule

2. Select **Rule Type** and select one of the following options:

    - *Antispyware*: Checks compliance to designated anti-spyware requirements.
    - *Antivirus*: Checks compliance to designated anti-virus requirements.
    - *CVE check*: Checks for protection against a list of publicly disclosed Common Vulnerability and Exposure (CVE) notices (Windows client devices only).
    - *Command*: Runs a command on the client device to check against an expected value (macOS client devices only).
    - *File*: Checks for the existence of a known file on the client.
    - *Firewall*: Checks compliance to designated firewall requirements.
    - *Hard Disk Encryption*: If encryption software is installed on the client device, this rule type checks the device's hard disks for applied encryption.
    - *Location*: Checks the client device's geographic location matches, or avoids, a list of defined locations.
    - *Mac Address*: Checks the client device's MAC address.
    - *Netbios*: Checks the client device's Netbios domain name.
    - *Network*: Checks the client device complies with a defined IP address and netmask range.
    - *OS*: Checks the client device’s Operating System meets a defined minimum standard.
    - *Process*: Checks for the existence of a known process on the client.
    - *Port*: Checks the client device's network interface ports.
    - *Patch Management*: If patch management software is installed on a client device, this rule type checks for the existence of missing software patches.
    - *Registry*: Checks for a value in a registry key (Windows client devices only).
    - *Risk Sense*: Supports Allow access, Block access and Notify based on the risk level.
    - *System Integrity*: Checks the system integrity of the client device (macOS client devices only).

- *Time of day*: Checks resource access requests against compliance with a time-based access schedule.

Restrictions exist for rule type availability on the following *Ivanti Secure Access Client* platform variants:
- Android clients are limited to rules based on *jail_break_root* and *OS*.
- iOS clients are limited to rules based on *jail_break_root*, *OS*, and *Time of day*.
- Linux clients are limited to rules based on *File*, *Port*, and *Process*.

3. Enter a **Rule Name** for your device rule.
4. (Optional) Enter a **Rule Description** for your device rule.

5. The remaining options are dependent on the **Rule Type** you selected:

   For *Antispyware* and *Firewall* rules, see Options for Antispyware and Firewall Rules.

   For *Antivirus* rules, see Options for Antivirus Rules.

   For *CVE check* rules, see Options for CVE Check Rules.

   For *Command* rules, see Options for Command Rules.

   For *File* rules, see Options for File Rules.

   For *Hard Disk Encryption* rules, see Options for Hard Disk Encryption Rules.

   For *Location* rules, see Options for Location Rules.

   For *Mac Address* rules, see Options for MAC Address Rules.

   For *Netbios* rules, see Options for Netbios Rules.

   For *Network* rules, see Options for Network Rules.

   For *OS* rules, see Options for OS Rules.

   For *Process* rules, see Options for Process Rules.

   For *Port* rules, see Options for Port Rules.

   For *Patch Management* rules, see Options for Patch Management Rules.

   For *Registry* rules, see Options for Registry Rules.

   For *Risk Sense* rules, see Options for Risk Sense Rules.

   For *System Integrity* rules, see Options for System Integrity Rules.

   For *Time of day* rules, see Options for Time of Day Rules.

6. Select **Create Rule** to create the device rule.

   The new rule is added to the list of device rules.

Individual device policies cannot be referenced by a secure access policy. After you have created all required device policies, you must organize them into device policy groups, see Creating Device Policy.

Editing / Deleting a Custom Device Policy Rule

To edit a device policy rule:

1. In the Device Policies page, under the Rules column, click the rule link that you want to modify.

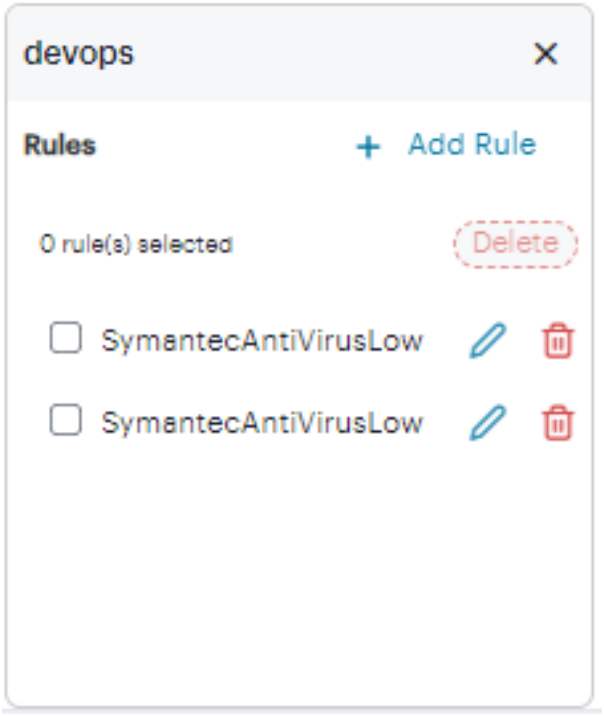

Edit Device Rule

2. In the side-panel that is displayed:

   - To add more rules, click **Add Rule**. To learn more, see "Creating Device Policy Rules" on page 451.
   - To delete a rule, click the delete icon next to the rule that you want to delete.
   - To delete more than one rule, select the check boxes next to the rules that you want to delete, and then click **Delete**.

Options for Antispyware and Firewall Rules

1. Select **Platform** and select one of the following options:

   - *windows*
   - *mac*

   Using the selected platform, *nZTA* populates the lists of *Vendors* and *Products* that can be selected for this rule.

2. (Optional) Select **Select Vendors** and use the drop-down list to select or deselect one or more product vendors. When done, select anywhere outside of the list.

   Each selected vendor is added to the panel below the drop-down list. To remove a selection, select the corresponding **X** indicator.

3. (Optional) Select **Select Products** and use the drop-down list to select or deselect one or more products. When done, select anywhere outside of the list.

   Each selected product is added to the panel below the drop-down list. To remove a selection, select the corresponding **X** indicator.

   

   While both *Vendor* and *Product* fields are optional, you must select at least one vendor or product for your rule.

4. (Optional) To set advanced options for this rule, select **Advanced Configuration**.

   The following options are provided:

   - Enable monitoring of this rule in *Ivanti Secure Access Client*.

Options for Antivirus Rules

1. Select **Platform** and select one of the following options:

   - *windows*
   - *mac*

   Using the selected platform, *nZTA* populates the lists of *Vendors* and *Products* that can be selected for this rule.

2. (Optional) Select **Select Vendors** and use the drop-down list to select or deselect one or more product vendors. When done, select anywhere outside of the list.

   Each selected vendor is added to the panel below the drop-down list. To remove a selection, select the corresponding **X** indicator.

3. (Optional) Select **Select Products** and use the drop-down list to select or deselect one or more products. When done, select anywhere outside of the list.

   Each selected product is added to the panel below the drop-down list. To remove a selection, select the corresponding **X** indicator.

   While both *Vendor* and *Product* fields are optional, you must select at least one vendor or product for your rule.

4. Select **Enforcement Level** and select one of the following options:

   - *high*
   - *moderate*
   - *low*

5. (Optional) To set advanced options for this rule, select **Advanced Configuration**.

   The following options are provided:

   - Add a maximum allowed time limit since the last successful system scan, in days.
   - Add a maximum allowed age limit for the most recent virus definition file update, either by number of available updates or by number of days.
   - Enable monitoring of this rule in *Ivanti Secure Access Client*.

Options for CVE Check Rules

This rule type is applicable to Windows devices only.

1. Select one of the following options:

   - To check all supported CVEs, select **Require all supported CVE checks**.
   - To check a list of specific CVEs, select **Check for specific CVE**, then use the **Select CVE Checks** drop-down control to select or deselect CVEs to be included.

To remove a selected CVE from the list, select the "X" button adjacent to the CVE tag.

Options for Command Rules

This rule type is applicable to macOS devices only.

In this release, Command Type is limited to "Defaults Read Command" only. This runs the `/usr/bin/defaults read` command on the client device.

1. Enter a value in **Argument1** to represent the path of the *Property List* file to read. For example, `/Applications/Utilities/Terminal.app/Contents/Info.plist`.
2. Enter a value in **Argument2** to represent the property key name. For example, `CFBundleShortVersionString`.
3. Enter one or more **Expected Values** to be returned by the command, as a comma-separated list. "*" (wildcard) values are also accepted.

Options for File Rules

This rule type is applicable to Windows and macOS devices only.

1. Select **Platform** and select one of the following options:
   - *windows*
   - *mac*
   - *linux*
2. Enter a full file name and path in **File Name**. For example, "c:test.txt" or "/Users/exampleuser/Downloads/test.txt".
3. Select **Checksum Type** and select one of the following options:
   - *md5*
   - *sha256*
4. Enter the **Checksum** value for the file.

5. Select **Mode** and select one of the following options:
    - *allow*. Select this to allow access where the file exists and is valid.
    - *deny*. Select this to deny access if the file does not exist or is invalid.

Options for Hard Disk Encryption Rules

This rule type is applicable to Windows and macOS devices only.

1. Select the device **Platform** to which this rule applies.
2. Select the **Vendors** and associated encryption **Products** you want this rule to check.
3. Choose which hard drives you want the rule to check:
    - To check all drives detected on the client device, select **All Drives**.
    - To check specific drives on the client device, select **Specific Drives**, then enter the drive identifiers required.
4. Select **Advanced Configuration** to provide additional rule configuration:
    - (*Specific drives* only) To ensure the rule does not trigger a failure where one or more of the specified drives are not detected, select **Consider policy as passed if the drives are not detected**.
    - To ensure the rule does not trigger a failure where detected drives are currently undergoing encryption, but are not yet fully encrypted, select **Consider policy as passed if the drive encryption is in progress**.

Options for Location Rules

1. Select **Mode** and select one of the following options:
    - *allow*. Select this to enable access for devices identified as being present at one of the set locations in the rule.
    - *deny*. Select this to disallow access for devices identified as being present at one of the set locations in the rule.

2. Use the "Add a location" section to define one or more geographic locations to which the current **Mode** applies:
    - Select a **Country**, **State** (optional), and **City** (optional).
    - To add the location, select **Add**.
3. Repeat the above steps for each location you want to add to the rule. Multiple "allow" and "deny" locations are possible in a single rule, with each added location identified by a green (allow) or red (deny) tag in the list.

    To remove a location, select the "X" button adjacent to the location tag.

Options for MAC Address Rules

1. Select **Platform** and select one of the following platform options:
    - *windows*
    - *mac*
2. Enter the **MAC address** as a comma-separated list (without spaces) of MAC addresses in the form HH:HH:HH:HH:HH:HH where the HH is a two-digit hexadecimal number. Duplicate MAC addresses are not supported.
3. Select **Mode** and select one of the following options:
    - *allow*. Select this to enable access from a listed MAC address.
    - *deny*. Select this to disallow access from a listed MAC address.

Options for Netbios Rules

1. Select **Platform** and select one of the following platform options:
    - *windows*
    - *mac*
2. Enter the Netbios domain **Names** as a comma-separated list (without spaces) of domain names. Each name can be 15 characters. Duplicate names are not supported.

3. Select **Mode** and select one of the following options:
    - *allow*. Select this to enable access from a listed Netbios domain name.
    - *deny*. Select this to disallow access from a listed Netbios domain name.

Options for Network Rules

1. Enter the **IP Address** and **Netmask** from which you want to either allow or deny access.
2. Select **Mode** and select one of the following options:
    - *allow*. Select this to enable access for the given IP address and netmask.
    - *deny*. Select this to disallow access for the given IP address and netmask.

Options for OS Rules

1. Select **Platform** and select one of the following options:
    - *windows*
    - *mac*
    - *ios*
    - *android*

2. The remaining fields are dependent on your choice of **Platform**:

    - Where you selected a platform of *windows* or *mac*, select **OS Name** and select an Operating System edition. For example, "Windows 2008" or "macOS Mojave".

      Then, select **OS Version** and select the version number or service pack associated with that edition of the Operating System. For example, "SP2" or "10.14.3". To not enforce the version number, select "Ignore".

    - Where you selected a platform of *ios* or *android*, select **Equality** and select one of the following options pertaining to how you want to enforce Operating System versions numbers:

        - *above*
        - *below*
        - *equal*

    Then, select **OS Version** and select the version number you want to check against.

Options for Process Rules

This rule type is applicable to desktop devices only.

1. Select **Platform** and select one of the following options:

    - *windows*
    - *mac*
    - *linux*

2. Enter a **Process Name**. For example, "explorer.exe".
3. Select **Checksum Type** and select one of the following options:

    - *md5*
    - *sha256*

4. Enter the **Checksum** value for the process executable.

5. Select **Mode** and select one of the following options:
    - *allow*. Select this to allow access where the process exists and is valid.
    - *deny*. Select this to deny access if the process does not exist or is invalid.

Options for Port Rules

1. Select **Mode** and select one of the following options:
2. Enter the **Ports** as a comma-separated list (without spaces) of ports. Port ranges are supported. Duplicate ports are not supported.
    - *windows*
    - *mac*
    - *linux*
3. Select **Platform** and select one of the following platform options:
    - *allow*. Select this to enable access from a listed port.
    - *deny*. Select this to disallow access from a listed port.

Options for Patch Management Rules

This rule type is applicable to Windows and macOS devices only.

1. Select the device **Platform** to which this rule applies.
2. Select the **Vendors** and associated patch management **Products** you want this rule to check the presence of.

3. (Optional) Select **Advanced Configuration** to view more options:

    - Choose the **Severity** levels of missing patches you want to check in this rule:

        - *Critical*
        - *Important*
        - *Moderate*
        - *Low*
        - *Unspecified/Unknown*

    For some products, the patch severity level might not be detectable. In this case, select *Unspecified/Unknown* to detect missing patches.

    - Choose the **Category** types of missing patches you want to check in this rule:

        - *Security Update*
        - *Rollup Update*
        - *Critical Update*
        - *Regular Update*
        - *Driver Update*
        - *Service Pack Update*
        - *Unknown*

    For some products, the patch category might not be detectable. In this case, select *Unknown* to detect missing patches.

Options for Registry Rules

This rule type is applicable to Windows devices only.

1. Select **Rootkey** and select one of the following options:
    - *HKEY_LOCAL_MACHINE*
    - *HKEY_USERS*
    - *HKEY_CURRENT_USER*
    - *HKEY_CURRENT_CONFIG*
    - *HKEY_CLASSES_ROOT*
2. Enter a **Subkey** for the registry path.
3. Select **Key Type** and select one of the following key types:
    - *string*
    - *dword*
    - *binary*
4. Enter a **Key** name.
5. Enter a **Value** for the registry key.
6. Tick the **64-bit** check box to use the 64-bit registry store. Leave this check box unticked to use the 32-bit registry store.

The following example values would create a rule to ensure the client device contained a registry key `HKEY_LOCAL_MACHINE\SOFTWARE\pzta` with a value `123`:

Field	Value
Rootkey	HKEY_LOCAL_MACHINE
Subkey	SOFTWARE
Key Type	string
Key	zta
Value	123
64-bit	*ticked*

Options for Risk Sense Rules

RiskSense provides vulnerability management and prioritization to measure and control cybersecurity risk. The cloud-based RiskSense platform uses a foundation of risk-based scoring, analytics to identify critical security weaknesses with corresponding remediation action plans, dramatically improving security and IT team efficiency and effectiveness.

Integrating RiskSense's Vulnerability Risk Rating (VRR) scores with nZTA provides an additional layer of security by isolating and preventing vulnerable devices from connecting to the nZTA network thereby protecting enterprise resources.

This rule type is applicable to Windows only.

1. Enter the **Rule Name**.
2. Enter the **Rule Details**.
3. Select **Risk Level** and select one of the following options:
    - Low
    - Medium
    - High
    - Critical
4. Select **Action** and select one of the following options:
    - Allow: Select this to allow access when the risk level is low or medium.
    - Block: Select this to block the access based on the risk level.
    - Notify: Select this to notify the user about the risk identified.

- RiskSense Alert will not be generated if the RiskSense device policy is enforced on the enrollment sign-in URL.
- RiskSense device policy should always be enforced on the authentication login URL.

To view Top Risky Applications, see the *Reviewing User Activity* section in Using the Insights Menu to Monitor User Activity and Service Usage.

Options for System Integrity Rules

This rule type is applicable to macOS devices only.

1. To enable this rule type, select **Enable**.

Options for Time of Day Rules

This rule type applies a resource restriction (allow or deny access) based upon a specified period frequency within a defined date and time range. Enter the following parameters:

1. Select the frequency with which you want the rule to apply inside the date range you specify:
    - **Custom**: Apply the rule for the whole period continuously between the start date/time and end date/time.
    - **Daily**: Apply the rule for the specified days in each month. Enter a comma-separated list of numerical days (1-31), for example: "1,5,19,28".
    - **Weekly**: Apply the rule for the specified days of each week. For **Select Days**, select the checkbox for each day on which you want the rule to apply.
    - **Monthly**: Apply the rule for all days in the specified months. For **Month**, select one or more months from the drop-down list.
2. Enter the **Start Date** and **End Date** to apply to the selected period frequency. For custom rules, the date range entered here is continuous. For daily, weekly, and monthly rules, each day in the range is executed individually according to the selected times and frequency.

    Start and end date values are optional for **Daily**, **Weekly**, and **Monthly** frequencies. If not specified, the rule applies indefinitely.

3. Enter the **Start Time** and **End Time** to apply to the selected period frequency. For custom rules, the times are applied with the corresponding start and end date to provide a continuous period within which the rule applies. For daily, weekly, and monthly rules, the times are applied for each day in the schedule.

   All times are applied as UTC timezone values. Your *nZTA Gateway*s must also use UTC time for the rule schedule to apply.

   Time periods for daily, weekly, and monthly rule frequencies are restricted to the 24 hours in a single day, such that you cannot enter an end time that is earlier than the start time. Therefore, in cases where you want to apply a rule allowing access for a time period that spans across midnight into the next day, add separate rules for each day in the range covering the time period for that day only. For example, to allow access during the period 21:00 Monday until 12:00 Tuesday, configure the following rules:
   Rule 1: **Period**: *weekly*, **Days**: *Monday*, **Start Time**: *21:00*, **End Time**: *23:59*, **Mode**: *allow*
   Rule 2: **Period**: *weekly*, **Days**: *Tuesday*, **Start Time**: *00:00*, **End Time**: *11:59*, **Mode**: *allow*

4. Choose the **Mode** that should apply during the specified times:

   - **allow**: Devices accessing resources to which this policy is applied are *authorized only* during the selected days and times.
   - **deny**: Devices accessing resources to which this policy is applied are *not authorized* during the selected days and times.

Setting Global Device Preferences

nZTA enables a system administrator to configure settings that control and restrict the functionality available in *Ivanti Secure Access Client* when a user enrolls their device with the *Controller*. Using the settings provided, you can control if your users are able to perform functions inside the *Ivanti Secure Access Client* application such as adding or removing connections, disconnecting from the *Controller*, or exiting the application completely.

Changes are replicated out to your end user devices at the point they next connect to the *Controller*.

To take advantage of the restriction settings described in this section, your users must be running the *Ivanti Secure Access Client* version applicable to *nZTA* 20.12 or later. To learn more about supported software versions, see the *Release Notes*.

These setting affect Windows and macOS desktop clients only. *Ivanti Secure Access Client* Linux variants are currently not supported.

To configure *Ivanti Secure Access Client* settings for your user's devices:

1. Log into the Tenant Admin Portal.
2. Click **Secure Access > Manage Devices**.
3. Click the **Global Device Preferences** tab.

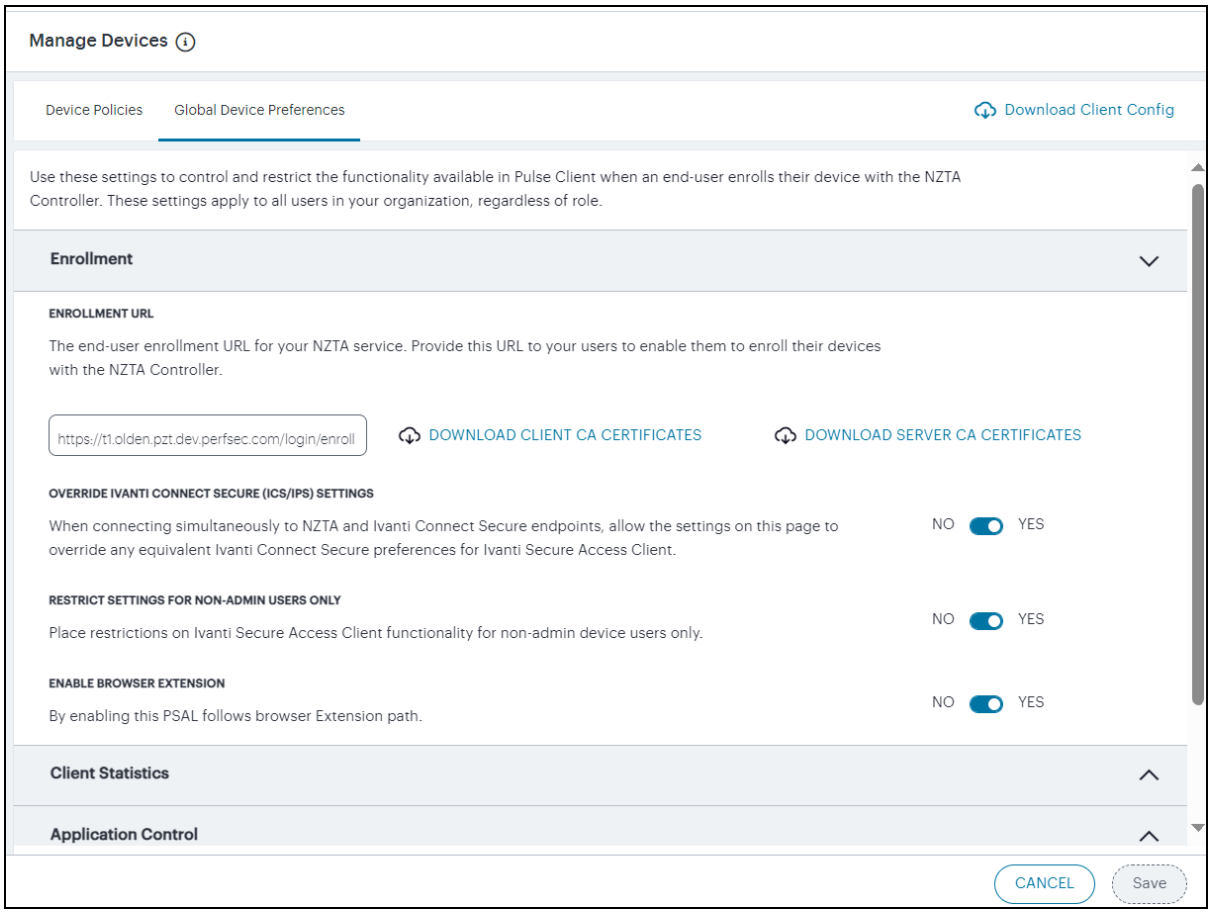

Configuring Global Device Preferences

Through this page, you can configure the following settings for *Ivanti Secure Access Client* on your end-user devices:

Ivanti Secure Access Client Settings

Setting	Category	Default Value	Description
Enrollment URL	Enrollment	None	A tenant-specific end-user enrollment URL. This setting is read-only, and can be used to inform your users of the correct *nZTA* enrollment URL.
Override Classic VPN (PCS/PPS) Settings	Enrollment	No	If your users use *Ivanti Secure Access Client* to simultaneously connect to classic VPN products from *Ivanti*, such as PCS or PPS, enable this setting to allow *nZTA* settings on this page to take precedence over any equivalent settings configured by the classic VPN. If you disable this option, *Ivanti Secure Access Client* functionality is determined by the classic VPN product you are connected to.
Restrict Settings for Non-Admin Users Only	Enrollment	No	By default, *Application Control* and *Connection Control* settings are enforced for all users. Enable this setting to apply the restrictions on this page to non-admin client device users only. Admin users are unaffected. For example, with this setting enabled, if **Allow DISCONNECT connection** is set to "No", a non-admin user is not allowed to disconnect a *nZTA* connection in the *Ivanti Secure Access Client* application whereas an admin user retains this capability.

Setting	Category	Default Value	Description
Start With Splash Screen	Application Control	Yes	Display the splash screen when launching the *Ivanti Secure Access Client* application.
Disallow Pulse Application Exit	Application Control	No	Prevent the end user from exiting the *Ivanti Secure Access Client* application.
Enable Embedded Browser	Application Control	Yes	This enables PSAL to follow browser extension path. Chrome/Edge browser to install and launch Ivanti Secure Access Client.
Suppress EUP Auto Launch	Application Control	No	Prevent the end user portal auto launch.
Allow Add New Connection	Connection control	Yes	Allow the end user to add new connections in *Ivanti Secure Access Client*.
Allow Delete Connection	Connection control	Yes	Allow the end user to delete connections in *Ivanti Secure Access Client*.
Allow Disconnect Connection	Connection Control	Yes	Allow the end user to disconnect a *nZTA* connection in *Ivanti Secure Access Client*.
Save User Credentials	Connection Control	No	By default, users cannot save and re-use their username and password credentials with a *nZTA* connection. Enable this setting to allow credentials to be saved.
Enable Always on Mode	Always on and Lock Down Mode	No	Always-on Mode allows the Ivanti Secure Access Client to establish a connection that is always active. The feature

Setting	Category	Default Value	Description
			restricts the users to manually connect/disconnect nZTA connection.
Enable Lock Down Mode	Always on and Lock Down Mode	No	If the tunnel is disconnected, for any reason, the machine has limited connectivity (only traffic allowed with exception rules) required to re-establish the tunnel. Always-ON mode with Lockdown mode enabled denies all network traffic until connected via nZTA connection. Exemption rules can be setup to allow network traffic.

Configuring Lock Down Mode

To enable Lock down this connection option, follow the below steps:

1. Select **Secure Access > Manage Devices > Global Device Preferences**.
2. Select **Enable Always ON mode** and **Enable LockDown Mode** option.
3. Click **View Exceptions**. When Always-on mode feature with Lockdown mode is enabled, Admin can add more exceptions to the Core Access Rules using exception rules. Exceptions already configured in the client are called Core Access Rules. DHCP, DNS, Kerberos, LDAP, SMP and Portmapper are already configured as Core Access Rules in the client. Exception rules can be configured to exempt certain types of traffic.

4. Click **Add** to add exception.

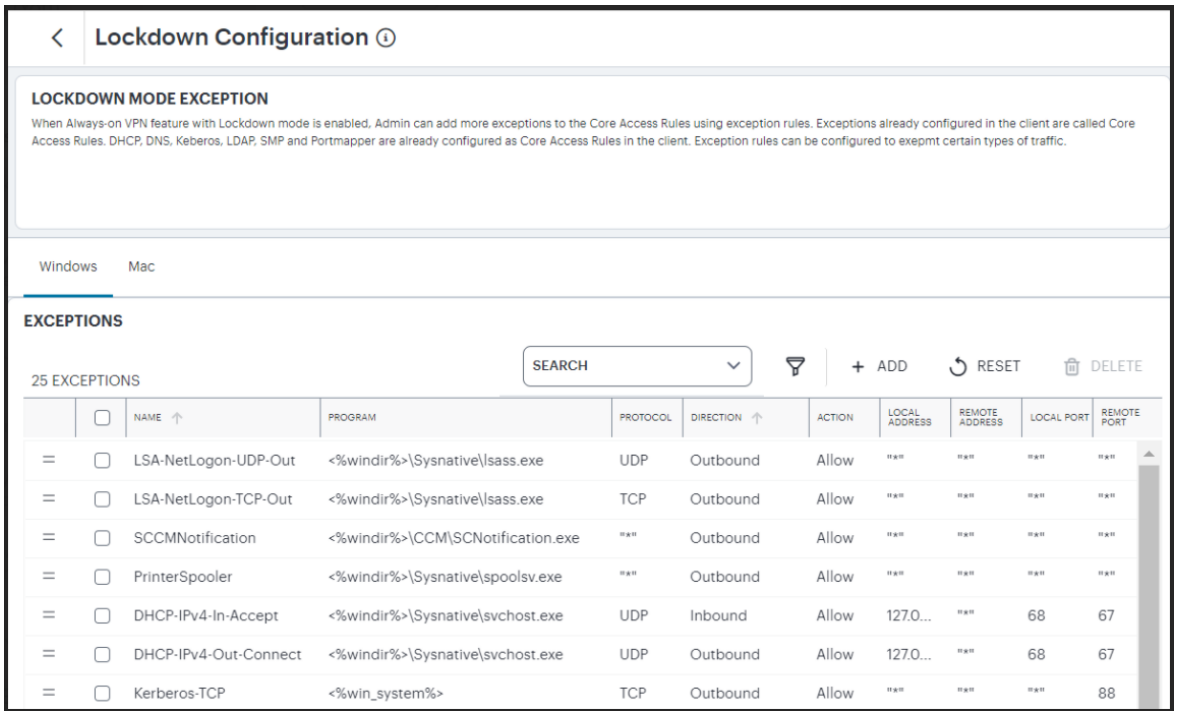

Lock Down configuration

5. Select the Platform **(Windows/Mac)**.
6. Enter the exception **Name** and **Description**.
7. Select the type:
    - Program
    - Port
    - Custom
8. Select the traffic type.
    - Inbound traffic is always directed towards user's machine.
    - Outbound traffic is always directed towards outside the machine.
    - Select Allow or Deny actions to configure the exception rules.
9. Click **Add Exception**.

Downloading Device Preferences for use with an External Service

nZTA provides the facility to download a file containing the device preferences and settings on the **Global Device Preferences** tab, in JSON format, for use with external Mobile Device Management (MDM) and Mobile Application Management (MAM) services, such as *Microsoft Intune* or *Jamf* (for Apple devices). This can be useful in enabling your end-users to receive the *Ivanti Secure Access Client* package from your MDM/MAM service along with preset configuration representing your specific enrollment details. Thus, your end-users need not perform post-installation configuration of *Ivanti Secure Access Client*, and can instead connect straight to the *nZTA* service.

To download the device preferences file, use the *Download* icon at the top of the page:

Downloading Global Device Preferences as a JSON file

This link activates a download dialog to save the device preferences file, in JSON format, to your local workstation.

To use the device preferences file with *Ivanti Secure Access Client*, base64-encode the JSON contents. This is necessary to enable the JSON structure to be presented as a singular string input argument to the *Ivanti Secure Access Client* package. Several freely-available applications such as text editors, or online services, provide this facility.

Then, to specify the configuration as a command-line argument to the *Ivanti Secure Access Client* package executable, copy and paste the encoded string into the **JSONCONFIG** argument using the following syntax:

```
JSONCONFIG="<base64-encoded-config>"
```

For example:

```
JSONCONFIG="eyJhcHBsaWNhdGlvbl9jb250cm9sIjp7ImRpc2FsbG93a...ddjsK
a435sag"
```

Browser-based interfaces such as Intune provide command-line argument specifiers as part of the application definition. Enter the complete `argument=value` string as shown above.

Providing the JSONCONFIG argument as part of the application definition on your end-user devices means that *Ivanti Secure Access Client* is installed fully-configured with the enrollment details provided by the downloaded device preferences file.

When using this mechanism to update an end-user device that has *Ivanti Secure Access Client* already installed, the *Ivanti Secure Access Client* software is upgraded as necessary, but pre-existing *nZTA* connections remain unaffected.

To learn more about how this facility might be used with your own MDM/MAM service, see your support representative.

Working with Applications and Application Groups

Introduction

After you have defined the user authentication system for your *Ivanti Neurons for Zero Trust Access* (*nZTA*) service, you can:

- Create definitions of applications to which your end users require access, see Adding Applications to the Controller.
- Group together multiple applications for which a single secure access policy is required, see Adding Application Groups to the Controller.
- Create secure access policies through which you can publish your applications, or application groups, to a *nZTA Gateway*, see Workflow: Publishing Applications to nZTA Gateways.

An application, or application group, can be associated with only one secure access policy.

Adding Applications to the *Controller*

For each application you want to make available through *nZTA*, you add an application definition to the *Controller*. Application definitions are referenced from a *secure access policy* in the following ways:

- A single application can be referenced from a secure access policy to identify an application for the policy.
- Multiple applications can be referenced from an application group, to enable all of the applications in the application group to be identified for a secure access policy.

To add an application:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Applications > Applications**.

   The *Applications* page appears. This page lists all applications defined on the *Controller*.

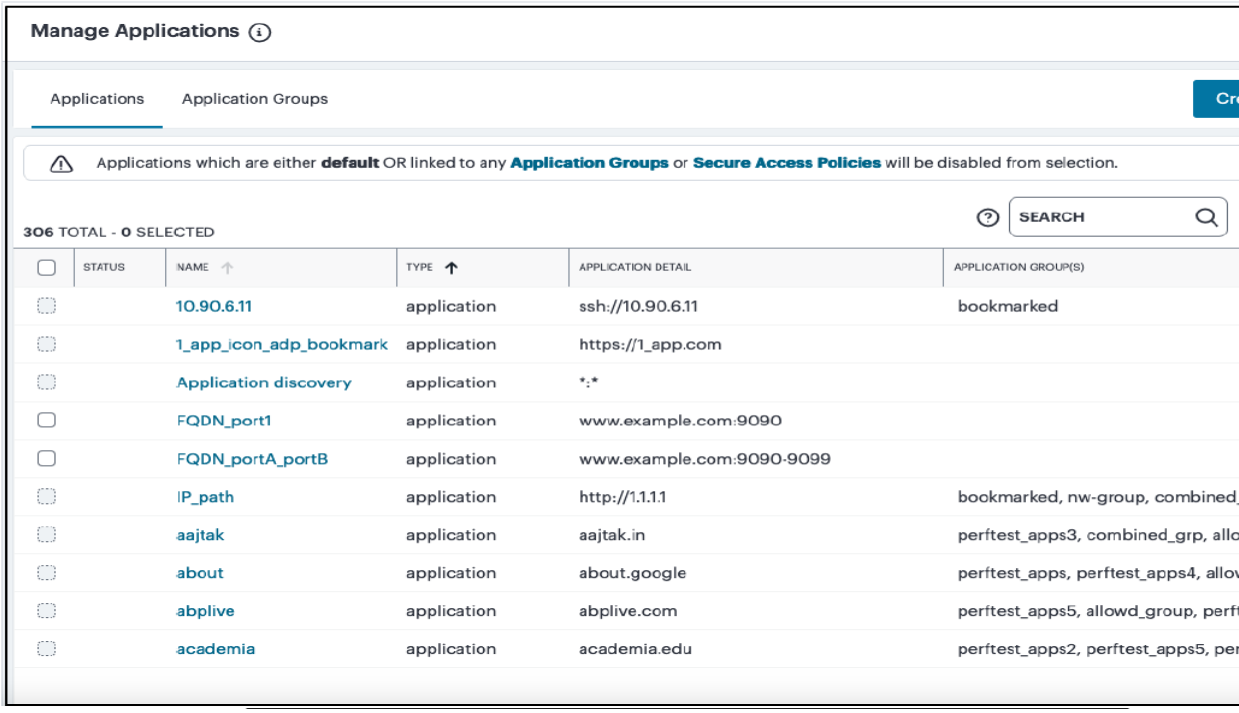

Manage Applications Page

This page also includes a built-in application called *Application discovery*. The **Application Detail** for this application is *:*, indicating that all applications that it applies to all unlisted applications. This application is used by the *nZTA* application discovery feature, and cannot be deleted.

3. Click **Create Application**.

   A form appears enabling you to create the application.

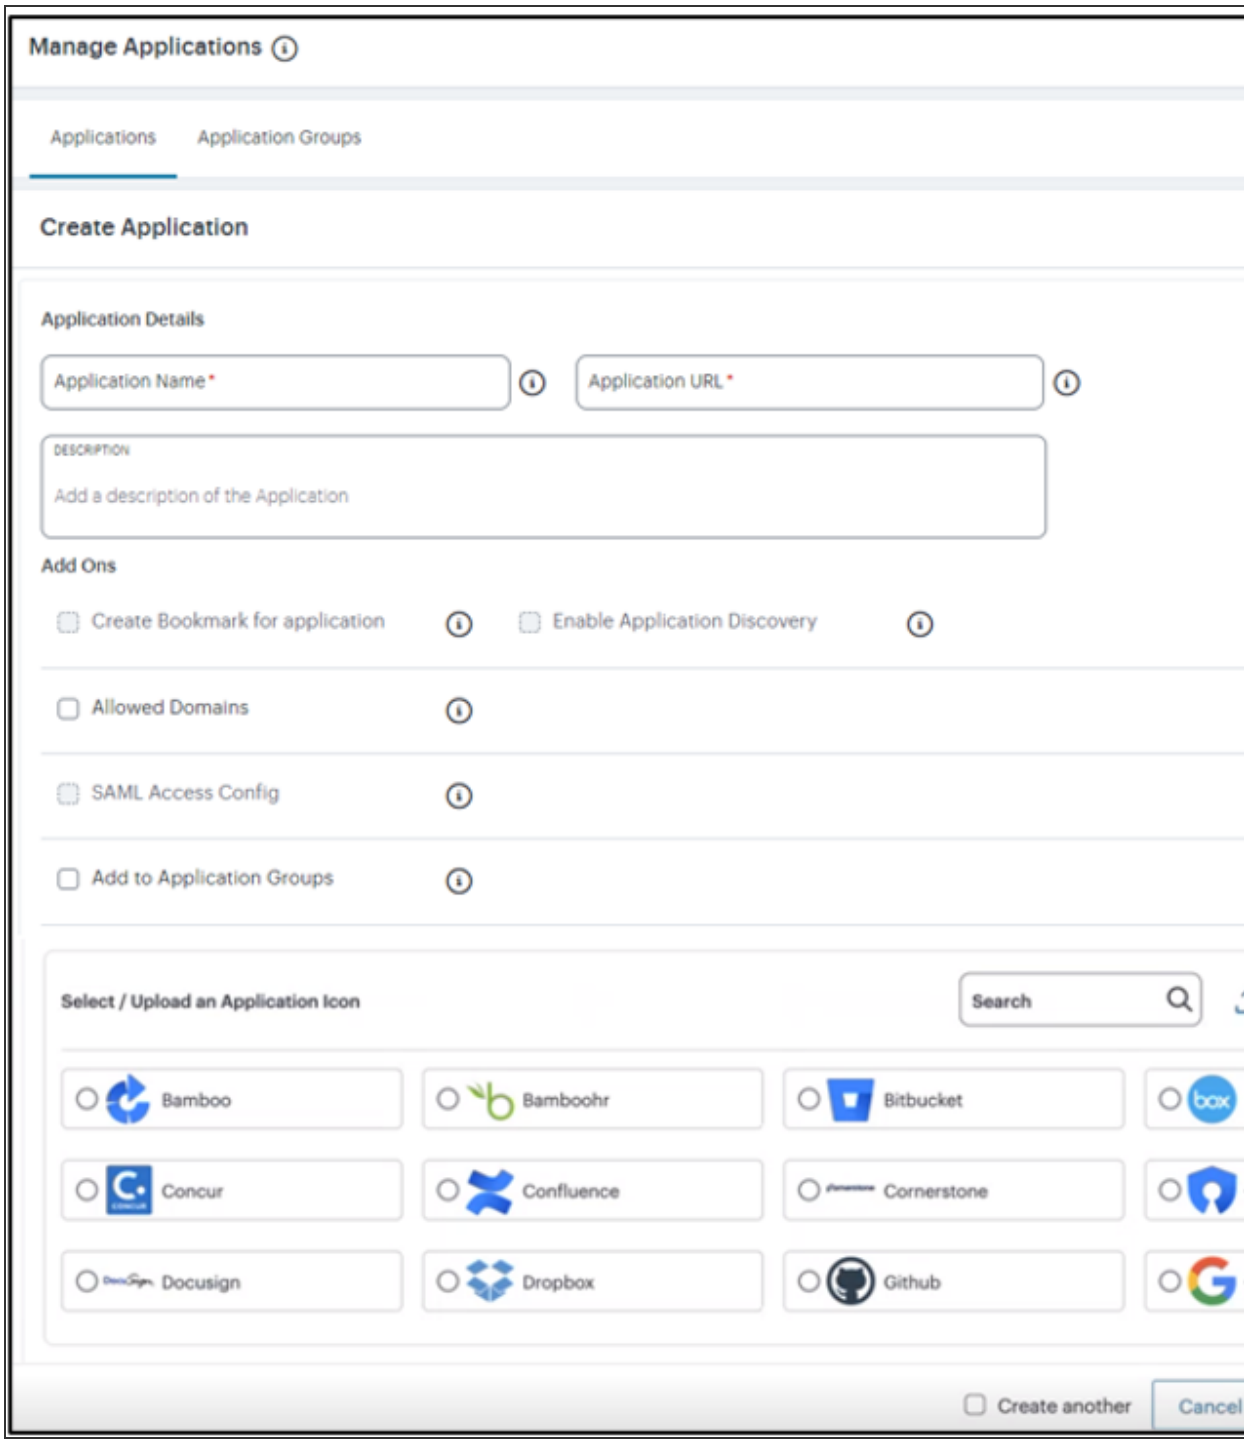

Create Application

4. Enter the **Application Name**.

5. Enter the **Application Details**. That is, the URI (Uniform Resource Identifier) you use to access the application. To view a complete list of valid entries for this field, see Defining Applications and Application Groups.

6. For scenarios that require one or more additional domains to be associated with an application, select **Allowed Domains**:

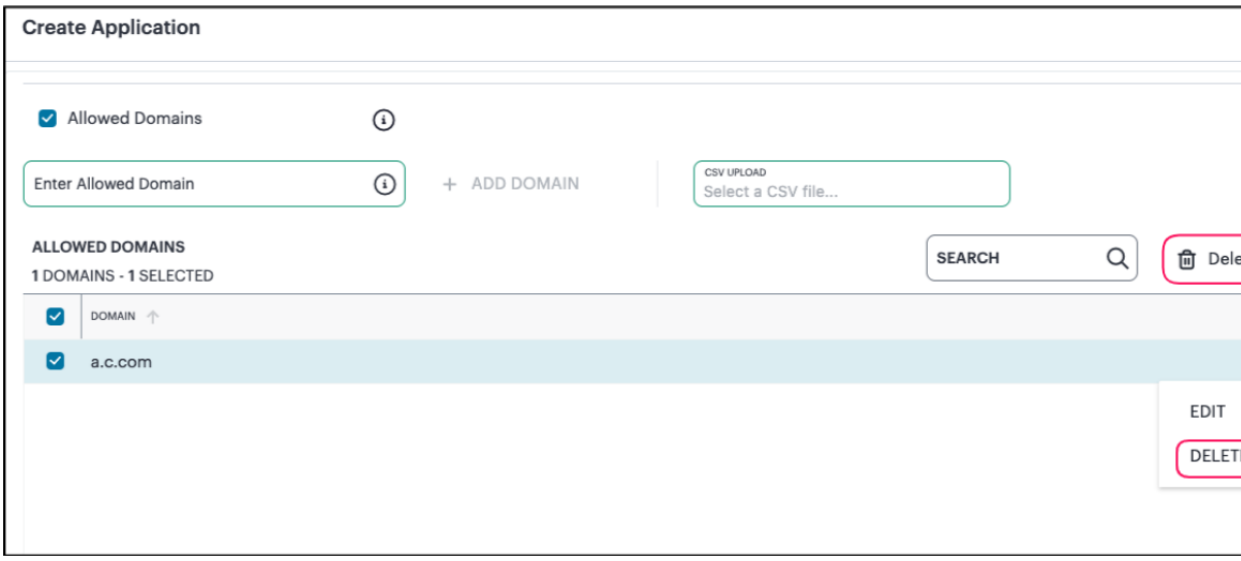

Adding allowed domains for an application

Add your domains through one of the following methods:

- Individually, by entering valid domains in the **Enter Allowed Domain** text box, then selecting **Add Domain** to add the domains to the list. You can add several domains at the same time by using a comma (,) separator. Repeat this step for each domain you want to add.
- In bulk, by uploading a Comma-Separated Value (CSV) text file containing the full list of your domains.

Domains added to this list must conform to the same scheme rules as the URI used in the **Application Details** field. To view a complete list of valid domain schemes, see Defining Applications and Application Groups.

In the list of added domains:

- to edit an entry, click the three dots next to the entry and then select **Edit**.
- to remove individual entries, click the three dots next to the entry and then select **Delete**.
- to remove all entries, select all check boxes and click **Delete**.

7. For HTTP/HTTPS applications, the **SAML Access** setting appears.

   The *Controller* can use SAML to provide a secure connection to your application or resource. In this scenario, *nZTA* acts as a SAML Identity Provider (IdP), with the application acting as the SAML Service Provider (SP). To learn more about using SAML, see SAML Authentication.

   - Disable this setting if you are using a application-level login for the application.
   - Enable this setting if you are using SAML single sign-on for the application. Then:
     - Under **Download IdP Metadata**, click **Download the IDP metadata file using the link** and save the IdP metadata file.
     - Log into your application and upload the IdP metadata file. Refer to the product documentation for the third-party application for details of this process.
     - In the application, download the SAML metadata as a file. Refer to the product documentation for the third-party application for details of this process.
     - Under **Upload the file below**, select and upload the SAML metadata file from the application.

     > You must keep the SAML metadata up-to-date, especially after renewing certificates. This is essential for a secure and successful SaaS Apps SAML SSO flow. Regularly updating configurations in both the Identity Provider and Service Providers helps prevent authentication failures and ensures the security of the authentication process.

8. (Optional) If you want to add custom SAML attributes, use **Attribute** and **Value** to add key-value pairs. Click **Add** to add an attribute pair, and repeat as required.

   Added attributes are displayed beneath the input fields. Click the corresponding **X** indicator to remove an attribute.

9. To associate an icon with this application, either:
    - Select an **Application Icon** from the list of supported icons. This field auto-populates based on the scheme you use in **Application Details**.
    - Click **Upload your own Icon** to upload a bespoke image file as the reusable custom icon. Then select the icon from the list to associate to this application. Make sure your icon is in JPEG format using the maximum dimensions 48 x 48 pixels (maximum file size 1 MB). *Ivanti* recommends you use only square images for your application icons. You can edit or remove the uploaded custom icon.
10. Enter a **Description** for the application.
11. (Optional) To create a bookmark for this application, select **Create bookmark for application**.

    Use the Bookmark option, where applicable, to allow the end user to copy the **Application Details** URI for use with other applications. For example, a TCP URI can be bookmarked to facilitate copy and paste into VNC or similar.

12. (Optional) To enable application discovery for this application, select **Enable Application Discovery**.

    To use application discovery, your application must be defined as a wildcard-prefixed FQDN (for example, "*.example.com"). To learn more about application discovery, see Defining Applications and Application Groups.

13. (Optional) If you want to add the new application to an application group, select the **Add to Application Group** check box, and then select the required application group.

    When using SAML authentication, make sure you add to a single application group only those applications that use the same SAML authentication source.

14. To save this application and create another application, select the **Create another** check box.

15. Click **Create Application**.

    The new application appears in the list of applications.

    Applications can also be added to the *Controller* during the **Create Secure Access Policy** workflow, see Workflow: Publishing Applications to nZTA Gateways.

    After you have defined your applications in the *Controller*, you can publish the actual applications to your *nZTA Gateways*, see Workflow: Publishing Applications to nZTA Gateways.

Editing and Deleting Applications

To edit an existing application definition, select the corresponding check box and click **Edit**. *nZTA* shows the *Edit Application* form, populated with the details of the application. Use this form to update the name and other details of your application.

For SAML applications, you can use the **Upload SAML Metadata** form to replace the metadata definition file previously-uploaded with a new or modified version. However, be aware that that federation metadata files can be digitally-signed and, in that case, cannot be manually edited prior to upload back into *nZTA*. In this scenario, you must obtain a new digitally-signed metadata file from your SAML SP suitable for uploading through this page. The parameters in an unsigned metadata file can be edited before the file is re-uploaded.

To delete an existing application, select the corresponding check box and click **Delete**.

You cannot delete the *Application discovery* application.

Adding Application Groups to the *Controller*

Multiple applications can be referenced from an *application group*.

When you select an application group during any subsequent process, all applications in the group are included automatically.

That is:

- During the creation of a secure access policy, see Creating/Editing Secure Access Policies.
- During the Create Secure Access Policy workflow, see Workflow: Publishing Applications to nZTA Gateways.

For SAML authentication, make sure you add to a single application group only those applications that use the same SAML authentication source. A secure access policy can associate an application group with only one authentication method. Therefore, all applications added to the group must use the same SAML metadata for authentication.

To create an application group:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Applications > Application Groups**.

   The *Applications Groups* page appears. This page lists all application groups defined on the *Controller*.

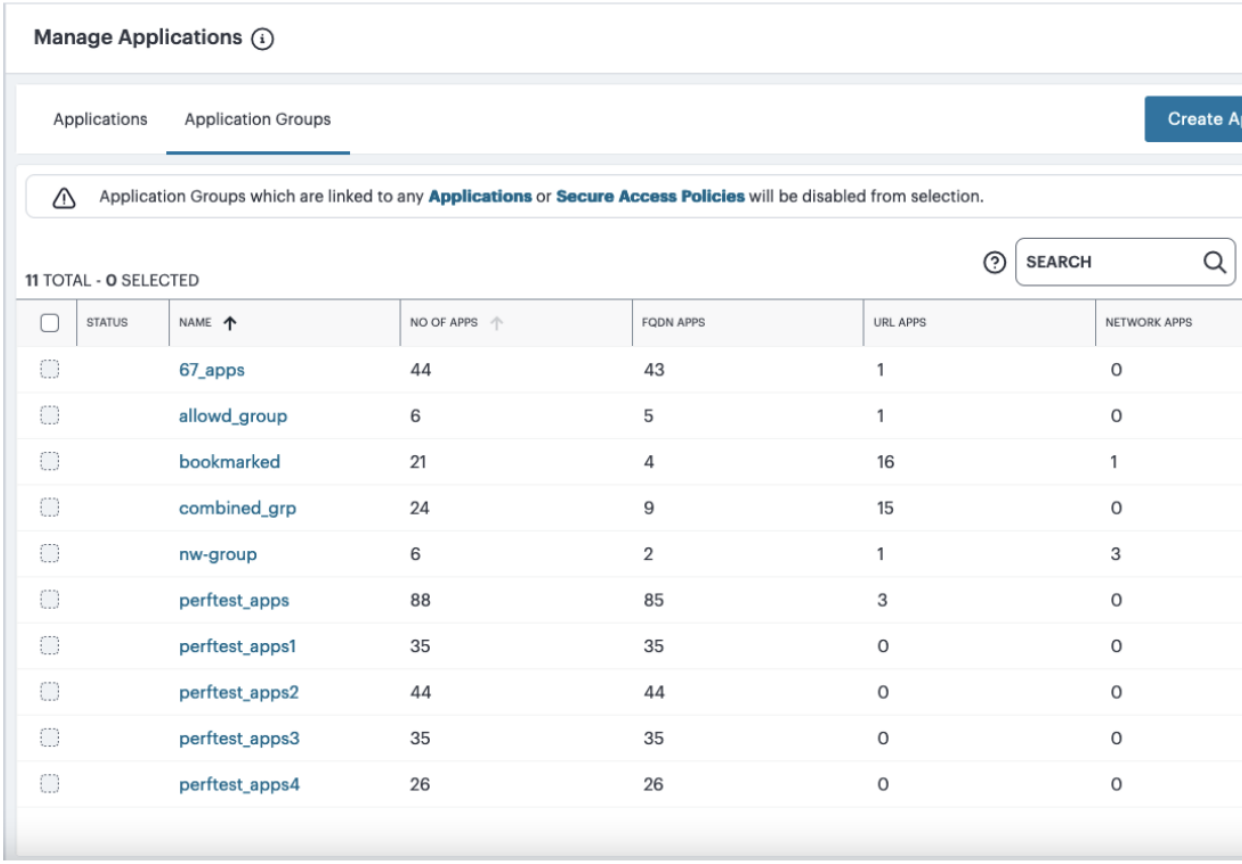

Manage Application Groups Page

3. Click **Create Application Group**.

   The **Create Application Group** form appears.

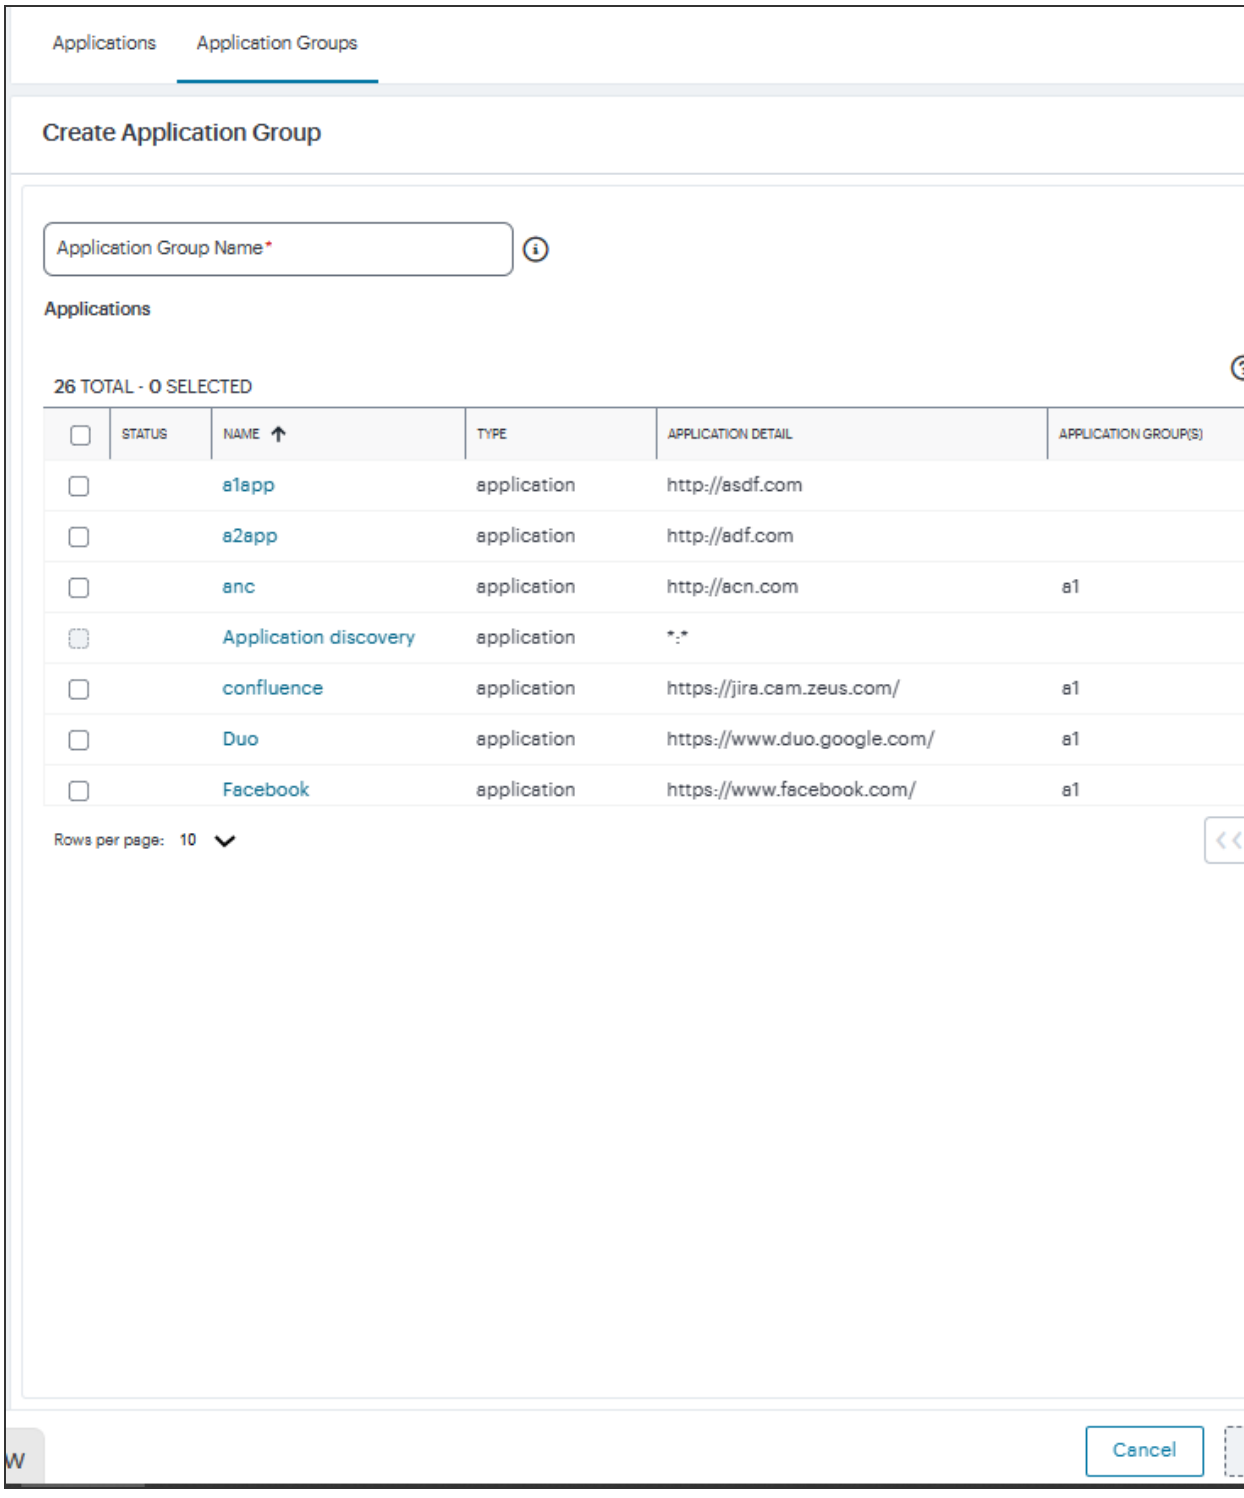

Create Application Group

4. Enter the **Group Name**.

5. Select the Applications you want to add to the group.

   You cannot add the *Application discovery* application to a group.

6. Click **Create Application Group**.

   The application group is added to the list.

Workflow: Publishing Applications to *nZTA Gateways*

After you have added any required application definitions to the *Controller*, you can publish these definitions to your *nZTA Gateway*(s) so that they are available for use.

To do this, use the **Create Secure Access Policy** workflow.

To publish applications to the *nZTA Gateway*(s), start the **Create Secure Access Policy** workflow.

You can access the **Create Secure Access Policy** workflow from:

- The **Secure Access Setup (Onboarding)** wizard, see Working with the Onboarding Wizard.
- The toolbar at the top-right of each page, see below.

To start the **Create Secure Access Policy** workflow using the toolbar:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.

   The *Network Overview* page appears.

2. Click the **Workflows** pull-down menu, and then select the **Create Secure Access Policy** workflow.

   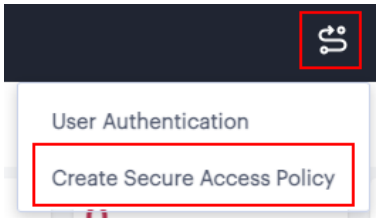

   Select the Create Secure Access Policy Workflow

   The **Create Secure Access Policy** workflow appears.

The **Create Secure Access Policy** workflow includes a multi-step workflow:

- **Select or Create Applications** that you want to publish, see Selecting Applications for Publication.
- **Select Device Policies** that apply to the application, see Selecting Device Policies for Applications.

- **Select or Create User Rules** that apply to the application, see Selecting User Rules for Applications.
- **Select Gateways** to which you want to publish applications, see Selecting a nZTA Gateway for your Applications.
- **Summary**, a confirmation-only step, see Confirming the Create Secure Access Policy Workflow.

After the **Create Secure Access Policy** workflow finishes, all selected applications are pushed to the selected *nZTA Gateway*.

If you are using multiple gateways, you will need to repeat the publication process for each gateway.

Selecting Applications for Publication

The **Select or Create Applications** step of the **Create Secure Access Policy** workflow enables you to create a new application, or to select an existing application that you want to publish.

You can also create applications independently of the Create Secure Access Policy workflow, see Adding Applications to the Controller.

To select an existing application:

1. Access the **Create Secure Access Policy** workflow, see Workflow: Publishing Applications to nZTA Gateways.
2. In the **Create Secure Access Policy** workflow, select the **Select or Create Application** step.
3. Click **Select an Application** and select the required application from the drop-down list.
4. (Optional) If you want to add the application to an application group, select the **Add to Group** check box, and then select the required application group.

   

   The applications in a group can be published as a single action.

   To learn more about the process of creating an application group, see Adding Application Groups to the Controller.
5. Click **Next** to continue to the next step of the workflow, see Selecting Device Policies for Applications.

To create a new application:

1. Access the **Create Secure Access Policy** workflow, see Workflow: Publishing Applications to nZTA Gateways.
2. In the **Create Secure Access Policy** workflow, select the **Select or Create Application** step.
3. Click **Select an Application** and select *Add New Application*.

   The add new Application form appears.
4. To add a new Application, follow the steps described in Adding Applications to the Controller.
5. (Optional) If you want to add the new application to an application group, select the **Add to Group** check box, and then select the required application group.

   The applications in a group can be published as a single action.

   To learn more about the process of creating an application group, see Adding Application Groups to the Controller.
6. Click **Next** to continue to the next step of the workflow, see Selecting Device Policies for Applications.

Selecting Device Policies for Applications

The **Select Device Policies** step of the **Create Secure Access Policy** workflow enables you to select the required device policy for the application that you want to publish.

To create device policies, see Creating Device Policies and Device Policy Rules.

To select device policies:

1. Access the **Create Secure Access Policy** workflow, see Workflow: Publishing Applications to nZTA Gateways.
2. In the **Create Secure Access Policy** workflow, select the **Select Device Policies** step.

   A list of existing device policies appears.
3. Select a device policy.
4. Click **Next** to continue to the next step of the workflow, see Selecting User Rules for Applications.

Selecting User Rules for Applications

The **Select or Create User Rules** step of the **Create Secure Access Policy** workflow enables you to compile a list of one or more user rules (and the groups to which they optionally belong) that apply to the applications you want to publish.

You can create user rules independently of the **Create Secure Access Policy** workflow, see Creating User Rules.

You can create user groups independently of the **Create Secure Access Policy** workflow, see Creating User Groups.

You can create authentication policies independently of the **Create Secure Access Policy** workflow, see Working with User Authentication.

To create a user rule:

1. Access the **Create Secure Access Policy** workflow, see Workflow: Publishing Applications to nZTA Gateways.
2. In the **Create Secure Access Policy** workflow, select the **Select or Create User Rules** step.
3. For the user group, either:
    - Click **Select or Create User Group(s)**, and select the required user group.
    - Click the *plus* symbol for the **Select or Create User Group(s)** property, and create the required user group using a **Group Name**, an **Authentication Policy** and (optionally) a **Description**.
4. For the authentication policy, either:
    - Click **Select an Authentication Policy**, and select the required policy.
    - Click the *plus* symbol for the **Select an Authentication Policy** property, and create the required authentication policy, see Working with User Authentication.

5. For the user rule, either:
    - Click **Select or Create Rule**, and select the required user rule.
    - Click the *plus* symbol for the **Select or Create Rule** property, and create the required user rule:
        - Enter a **Rule Name** for the rule.
        - Click **Select Attribute Type** and select the required authentication attribute type. The following options are supported: *Username*, *SAML (Azure AD)* and *Custom*.
        - Click **Expression** and select either *Matching* or *Not Matching*.
        - Enter the required **User** match string for the selected Expression. Wildcard matches are supported. For example: *
        - Click **Add to List**.
6. Click **Add User Rule**.

    The new user rule is added to the list of rules.
7. (Optional) Repeat steps 3 to 6 to create additional rules, if required.
8. In the list of rules, select each rule that is required by enabling its check box.
9. Click **Next** to continue to the final step of the workflow, see .Confirming the Create Secure Access Policy Workflow

To select an existing user rule:

1. Access the **Create Secure Access Policy** workflow, see Workflow: Publishing Applications to nZTA Gateways.
2. In the **Create Secure Access Policy** workflow, select the **Select or Create User Rules** step.

    The **Select or Create User Rules** page lists all existing user rules.
3. In the list of rules, select each rule that is required by enabling its check box.
4. Click **Next** to continue to the next step of the workflow, see Confirming the Create Secure Access Policy Workflow.

Selecting an *nZTA Gateway* for your Applications

The **Select Gateways** step of the **Create Secure Access Policy** workflow enables you to identify the *nZTA Gateway* to which you want to publish applications.

To select the required *nZTA Gateway*(s):

1. Access the **Create Secure Access Policy** page, see Workflow: Publishing Applications to nZTA Gateways.
2. On the **Create Secure Access Policy** page, select the **Select Gateways** step.
3. Click **Select Gateway** and select the required *nZTA Gateway*.
4. Click **Next** to continue to the next step of the workflow, see Confirming the Create Secure Access Policy Workflow.

Confirming the Create Secure Access Policy Workflow

After you have successfully completed all steps of the **Create Secure Access Policy** workflow, the final **Summary** step of the workflow becomes active.

This step displays all information that was defined/gathered during the **Create Secure Access Policy** workflow, and enables you to complete the workflow.

1. Access the **Create Secure Access Policy** workflow.
2. In the **Create Secure Access Policy** workflow, select the **Summary** step.

   A summary page displays all information that was defined/gathered during the previous steps.
3. Examine the summary information.
4. Click **Finish** to confirm the summary and complete the **Create Secure Access Policy** workflow.

   The applications are published to the selected *nZTA Gateway*.

After you have published applications to your *nZTA Gateway*(s), users can enroll their desktop and mobile devices, see Enrolling Mobile/Desktop Clients.

Creating/Editing Secure Access Policies

Introduction

A secure access policy defines how end-users can connect to *Ivanti Neurons for Zero Trust Access* (*nZTA*) to access applications. Each secure access policy is defined in terms of four dimensions:

- **Gateways**: An *nZTA Gateway* (or *nZTA Gateways*) through which Application access is controlled, and on which the secure access policy is deployed, see Working with Gateways.
- **Users**: The user methods, user policies, and user groups that are required on the *nZTA Gateway*, see Working with User Authentication.
- **Devices**: The device policy rules and device policies that define which end user devices can access applications, see Creating Device Policies and Device Policy Rules.
- **Applications**: The on-premise and cloud applications to which end users have access, see Working with Applications and Application Groups.

When a secure policy is created on the *Controller*, it is downloaded automatically and applied to the *nZTA Gateway* referenced by the policy.

nZTA has one built-in secure access policy, *Application discovery*. This policy, when enabled and configured, directs any request from an application that is not referenced by a policy to a *default Gateway*. See Configuring a Default Gateway for Application Discovery.

Viewing your Secure Access Policies

To see the list of secure access policies currently defined on the *Controller*:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, select the **Secure Access** icon, then select **Secure Access Policies**.

   The *Secure Access Policies* page appears.

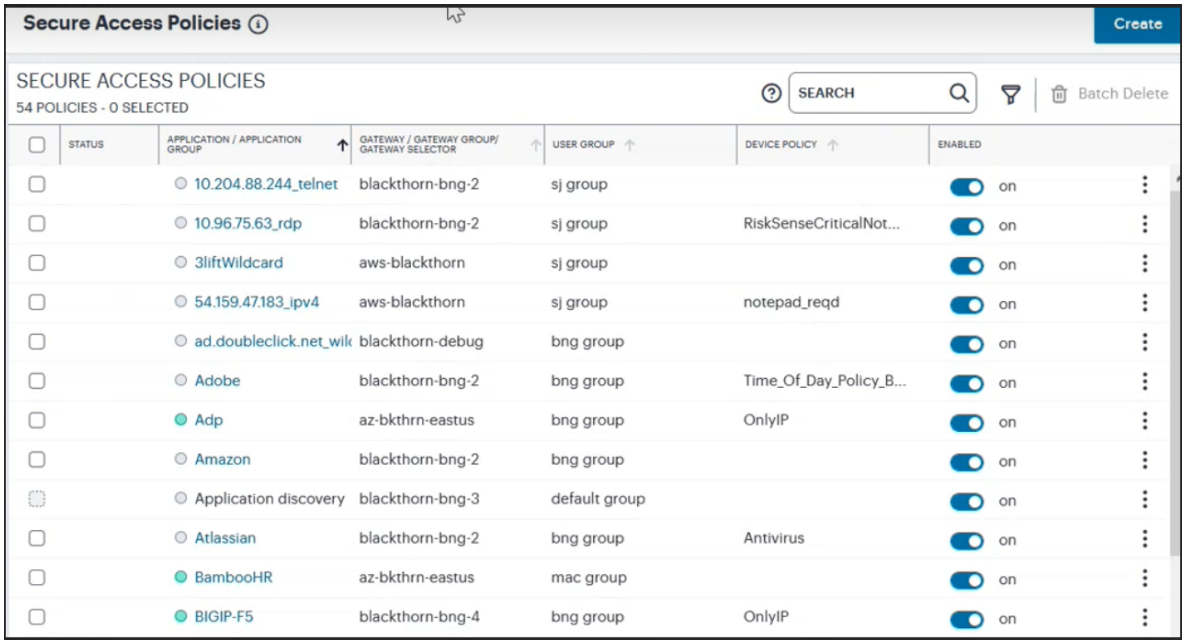

   Viewing the list of secure access policies

Each policy is shown with the following details:

- **Statusindicator**: The current availability of the application.
- **Applications/Application Groups**: The Application or Application Group this secure access policy is configured for.
- **Gateways**: The *nZTA Gateway* or *nZTA Gateway* Group or *nZTA Gateway* Selector that controls access to the applications named in this secure access policy.
- **Users**: The User Group for this secure access policy.
- **Device Policies**: The device policy that applies to this secure access policy.
- **Enabled**: Whether or not this secure access policy is enabled.

Select the application name to view the *Application* or *Application Groups* page.

An *application group* is identified by the arrow adjacent to the application name. Select the arrow to show the applications available in this group:

The applications contained in an application group

The *Status* indicator uses one of the following color schemes to determine application health:

- **Green**: The application (or application group) is reachable.
- **Gray**: The *Controller* cannot determine the availability of the application (or application group).
- **Red**: The application (or application group) is unreachable.

Application health status is available only for URL or IP address-based applications. Application definitions using a URI that contain a wildcard-based FQDN are not monitored for this purpose and are shown with a *gray* status indicator to denote that the status cannot be determined.

The built-in *Application Discovery* policy is not monitored for this purpose and uses a gray status indicator.

To view a chart showing the top non-reachable applications, see All Applications.

Through this page you can:

- *Create* a new secure access policy, see Creating a Secure Access Policy.
- *Edit* an existing secure access policy, see Editing a Secure Access Policy.
- Use the *View Details* link in the message, which appears on successful creating/updating of policy, to view the newly created/updated policy in the list.
- *Enable* or *disable* a secure access policy. Select the policies you want to enable/disable using the toggle button, then select the corresponding link at the top of the page.

- *Delete* a secure access policy. Use the three dots adjacent to the policies you want to delete, then select **Delete**.
- Perform *Search* for occurrences of named applications, application groups, gateways, gateway groups, device policies, user groups and enabled status (yes/no) for the policies listed on this page.
- *Filter* the policies displayed on the page by application/application group, gateway, user, device policy, or status. When you select the *Filter* icon, a side panel dialog appears within which you can select specific criteria to filter the display to show only matching policies. Applied filters remain in place until you select **Clear All** from the side-panel, or until you leave the page.

Creating a Secure Access Policy

A secure access policy defines how end users can connect to *nZTA* to access applications.

To create a secure access policy:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, select **Secure Access > Secure Access Policies**.

   The **Secure Access Policies** page appears. This lists all current secure access policies.

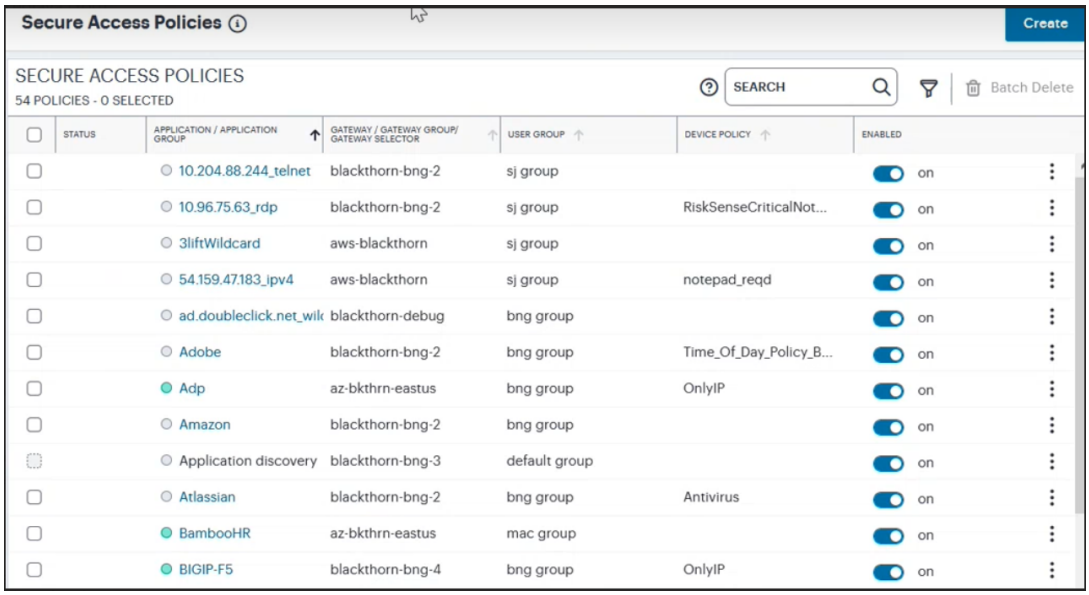

Viewing the Secure Access Policies page

3. Click **Create**:

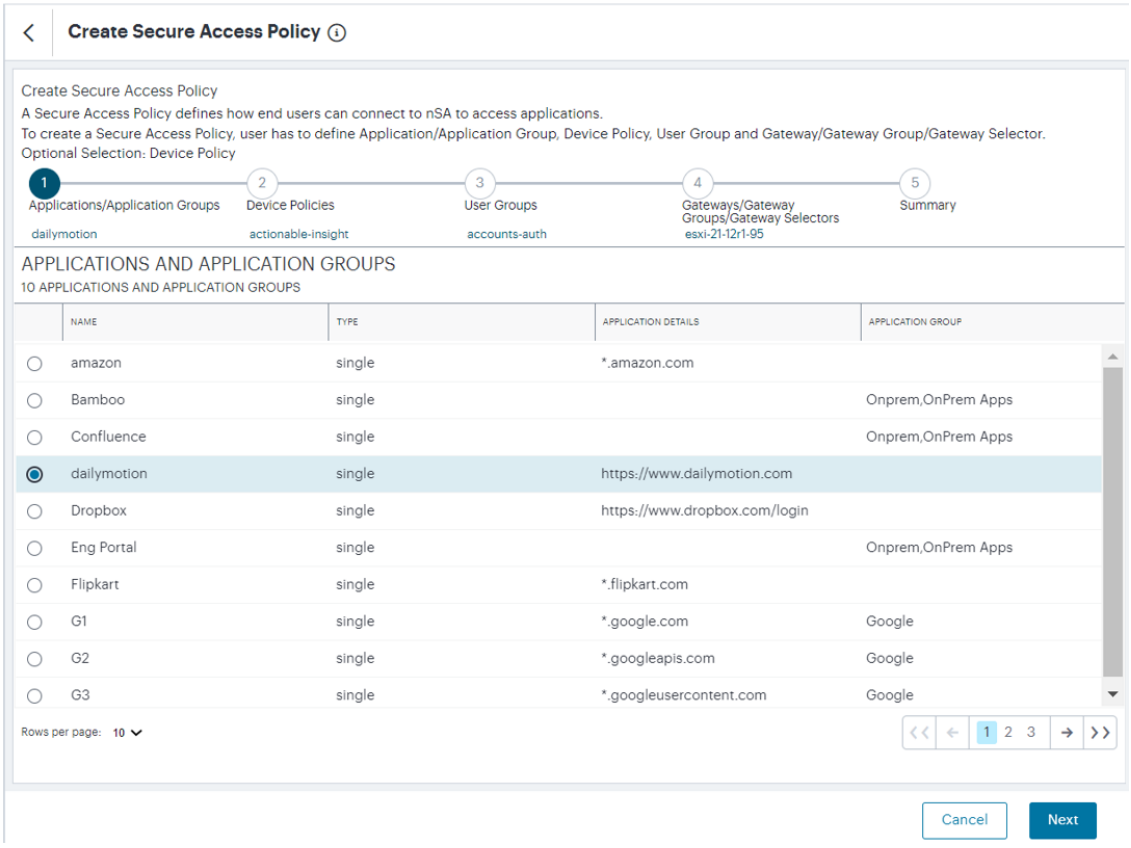

Creating a new Secure Access Policy

> At any point during this process, you can reset the form data by selecting **Reset**. You can also view existing secure access policies in a pop-up dialog by selecting **View Secure Access Policies**.

4. Select the application, or application group, for the policy. Click **Next**.

> An application, or application group, can be associated with only one secure access policy.

5. From the Device Policies list, select the device policy to apply to your Secure Access Policy. Click **Next**.
6. From the User Groups list, select the user group to apply to your Secure Access Policy. Click **Next**.

7. From the Gateways, Gateway Groups and Gateway Selectors list, select the nZTA Gateway/Gateway Group to which you want to publish your Secure Access Policy. Click **Next**.
8. Verify the Summary details and then click **Create**.

   The policy is created and added to the list of secure access policies.
9. (Optional) To edit a listed secure access policy, select the adjacent three dots and then select **Edit**. After the secure access policy is updated, it is automatically applied to the *nZTA Gateway* that it references.
10. (Optional) To enable a disabled secure access policy, use the toggle button. After the secure access policy is enabled, it is automatically applied to the *nZTA Gateway* that it references.
11. (Optional) To disable an enabled secure access policy, use the toggle button.
12. (Optional) To delete an *unused* secure access policy, select the adjacent three dots and then select **Delete**. Confirm the deletion in the subsequent dialog.

After the secure access policy is created, it is automatically downloaded and applied to the *nZTA Gateway* that it references.

Secure Access Policies can take several minutes to reach their destination *nZTA Gateway*(s). If an entered policy contains configuration that fails to apply properly due to a compatibility or validation problem, *nZTA* displays an error message that the applied configuration is incorrect. *nZTA* attempts to re-apply the configuration in the policy at 15 minute intervals, and repeats this process until such a time as the policy is corrected or deleted.

Editing a Secure Access Policy

You can edit a secure access policy to change how end users connect to *nZTA* to access applications.

For example, you must edit the *Application discovery* secure access policy if you want to use a default gateway to support the application discovery feature. To do this, you must assign an unused gateway to the secure access policy. See Configuring a Default Gateway for Application Discovery for details of default gateways.

To edit a secure access policy:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, select the **Secure Access** icon, then select **Secure Access Policies**.

   The *Secure Access Policies* page appears. This lists all current secure access policies.
3. Select the three dots adjacent to the secure access policy you want to edit, then select **Edit**.

   The **Edit Secure Access Policy** form appears.

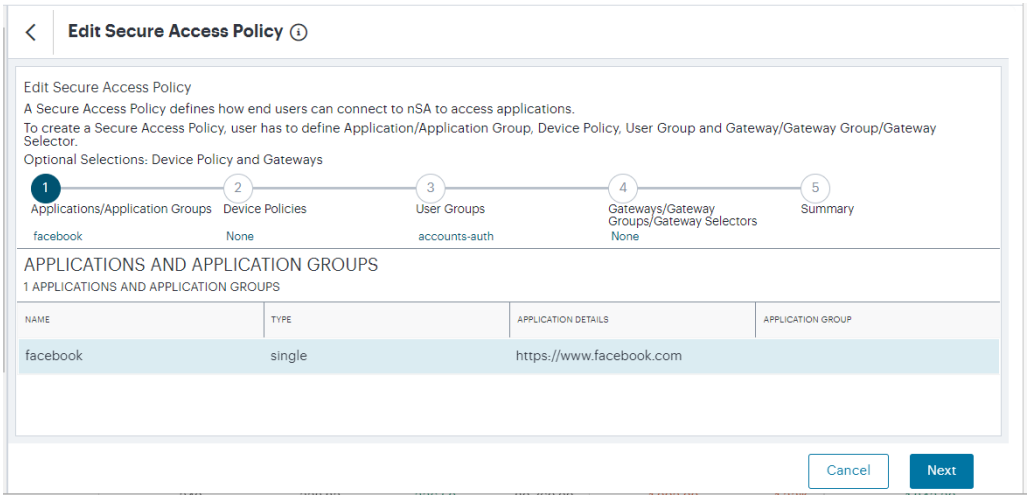

Editing a Secure Access Policy

> The **Application Type** and **Application** cannot be changed.

4. Make any required changes to the secure access policy.
5. Select **Save**.

Enrolling Ivanti Secure Access Client

Introduction

After you have created the required configuration for your *Ivanti Neurons for Zero Trust Access* (*nZTA*) service, you can begin to enroll your end user devices.

To see which devices have been enrolled, and to perform certain actions on enrolled devices, use the **Insights** > **Devices** page. For more details, see "Viewing Currently Enrolled User Devices" on page 666.

> For security reasons, only the authorized user account used to enroll a device is subsequently permitted to sign-in to *nZTA* on that device.

The following desktop and mobile device types are supported:

- *Ivanti Secure Access Client* on Windows OS desktop, see Enrolling a Windows Desktop Device.
- *Ivanti Secure Access Client* on macOS desktop, see Enrolling a macOS Desktop Device.
- *Ivanti Secure Access Client* on Linux desktop, see Enrolling a Linux Desktop Device.
- *Ivanti Secure Access Client* on iOS mobile device, see Enrolling an iOS Mobile Device.
- *Ivanti Secure Access Client* on Android mobile device, see Enrolling an Android Mobile Device.

The table below summarizes feature support for each device type:

Feature Support for Clients

Feature	iOS	Android	macOS	Windows	Linux
On-Demand	Yes	Yes	Yes	Yes	No

Feature	iOS	Android	macOS	Windows	Linux
Connection					
End User Portal	No	No	Yes	Yes	Yes (SSO,RDP,SSH not supported)
SAML Auto Sign-in & Single Logout (SLO)	No	No	Yes	Yes	Yes
Simultaneous Connection	No	No	Yes	Yes	No
Automatic *Ivanti Secure Access Client* Upgrade	Yes	Yes	Yes	Yes	No
Browser-Based Enrollment	No	No	Yes	Yes	No
Dynamic Policy/CARTA	No	No	Yes	Yes	Yes (CARTA message is not supported)
Device Policy	Yes	Yes	Yes	Yes	Yes (limited to support for File, Port number, and Process policy types)
FQDN/IP based Application	Yes	Yes	Yes	Yes	Yes (Split DNS is not supported)
Client Disconnect	Yes	Yes	Yes	Yes	Yes
Log Upload	Yes	Yes	Yes	Yes	No
nZTA Client Settings	No	No	Yes	Yes	No
CNAME FQDN App	Yes	Yes	Yes	Yes	No
Application	Yes	Yes	Yes	Yes	No

Feature	iOS	Android	macOS	Windows	Linux
Discovery Default Gateway					
GSLB Support	Yes	Yes	Yes	Yes	No
Customer PKI support (BYOC)	No	No	Yes	Yes	No
Multiple sign-in URL support	Yes	Yes	Yes	Yes	Yes

After a device is enrolled with *nZTA*, requests from each application are handled by the Gateway referenced in the secure access policy for the application.

A default Gateway can also be configured on the *Controller*, which handles requests from all applications that are not referenced by a secure access policy, see Application Discovery with Ivanti Secure Access Client.

Enrolling a Windows Device

Before you start this process, you must have an Windows sign-in URL for *nZTA*, based on the tenant FQDN provided by the *Ivanti* DevOps/Support organization.

If you have an existing *Ivanti Secure Access Client* installed, you must first uninstall it before beginning the *nZTA* enrollment process.

The Domain Admin must also ensure that Windows desktop machines can successfully connect to the Windows domain by updating the Trusted Root CA Certificates on all machines, see Enabling Trusted Root CA Certificate on Windows Domain.

To enroll a Windows desktop device:

1. Log into your Windows desktop.
2. Start a browser session.

3. In your browser, enter the *nZTA* enrollment URL into your address bar. For example, "https://tenant1.mycompany.com/login/" or "https://tenant1.mycompany.com/login/saleslogin/".

   A login page appears.

4. Provide your credentials to access *nZTA* enrollment.

   The launcher page appears, for example:

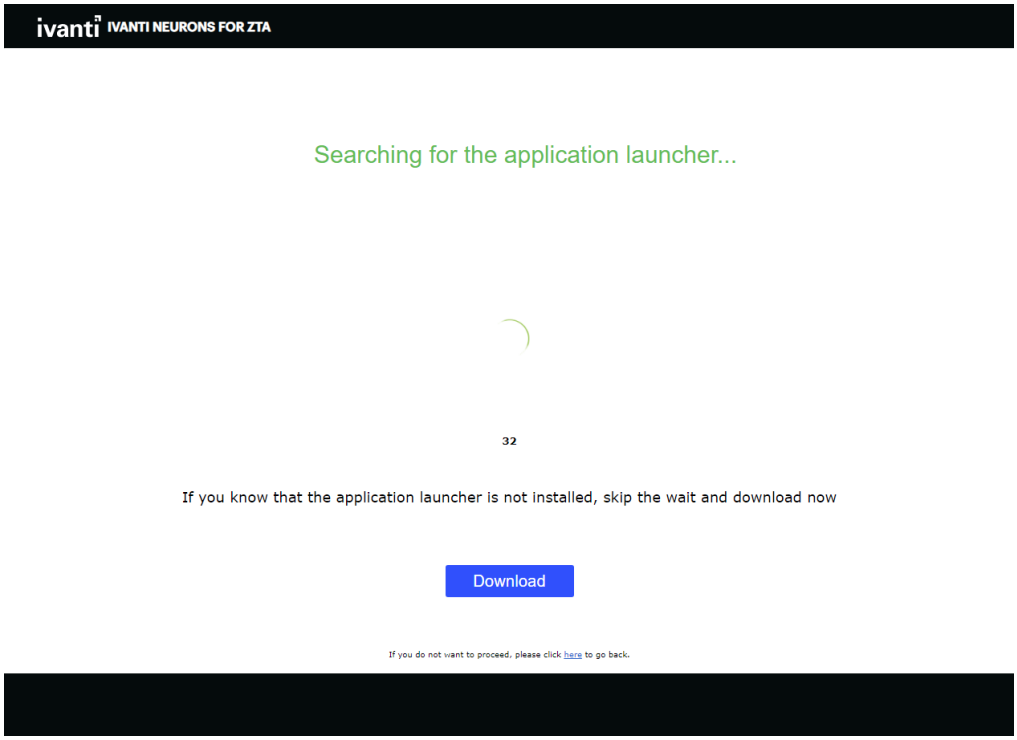

   nZTA Launcher (Windows)

5. Click **Download** and save the application launcher *PulseSecureAppLauncher.msi* file locally.

   > A "HERE" hyperlink is displayed in the browser. This is required *after* the launcher has installed.

6. Run the downloaded file to install the launcher.
7. Confirm the completion of the launcher installation.
8. In the browser, click the "HERE" hyperlink to continue with the enrollment.

   A browser dialog requests confirmation to start the launcher app.

9. Click **Open Pulse Secure Application Launcher**.

   The launcher starts.

10. Accept any certificate warnings.

    A progress bar indicates installation status.

11. Accept that the *PulseSetupClientOCX.exe* file can make changes to your device.

12. Accept that the *PulseSetupClientOCX64.exe* file can make changes to your device.

13. Accept that the *Pulse Secure Component Manager* Installer application can make changes to your device.

    Ivanti Secure Access Client then downloads.

14. Accept that the *Pulse Secure Component Manager* can make changes to your device.

    Ivanti Secure Access Client then installs and starts, and reports on a number of tasks:

    - *Enrolling the User*.
    - *Fetching and Importing Client Certificates*. You must confirm any certificate requests.
    - *Fetching and Importing CA Certificates*. You must confirm any certificate requests.
    - *Launching the Windows Edge/Webview2 browser*. In a typical enrollment, upon successful authentication to the Controller, *Ivanti Secure Access Client* automatically shows the end-user portal applications page through a Windows Edge/Webview2 browser. This feature is supported with ISAC client version 22.6R1.

    Onboarding is then complete.

    Ivanti Secure Access Client appears as a task in the Windows task bar. For example:

    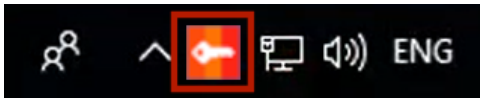

    Ivanti Secure Access Client Task in Taskbar (Windows)

15. When the *Controller* requests a certificate from the client, accept the request.

16. Log into the *Controller* using your *Ivanti Neurons for Zero Trust Access* service user credentials.

    The compliance of the device is checked.

17. (Optional) If your sign-in authentication policy is configured for Multi-Factor Authentication, you might be required to complete a TOTP (Time-based One Time Password) secondary authentication step each time you sign in.

    For a first time login, the user is presented with a TOTP registration page:

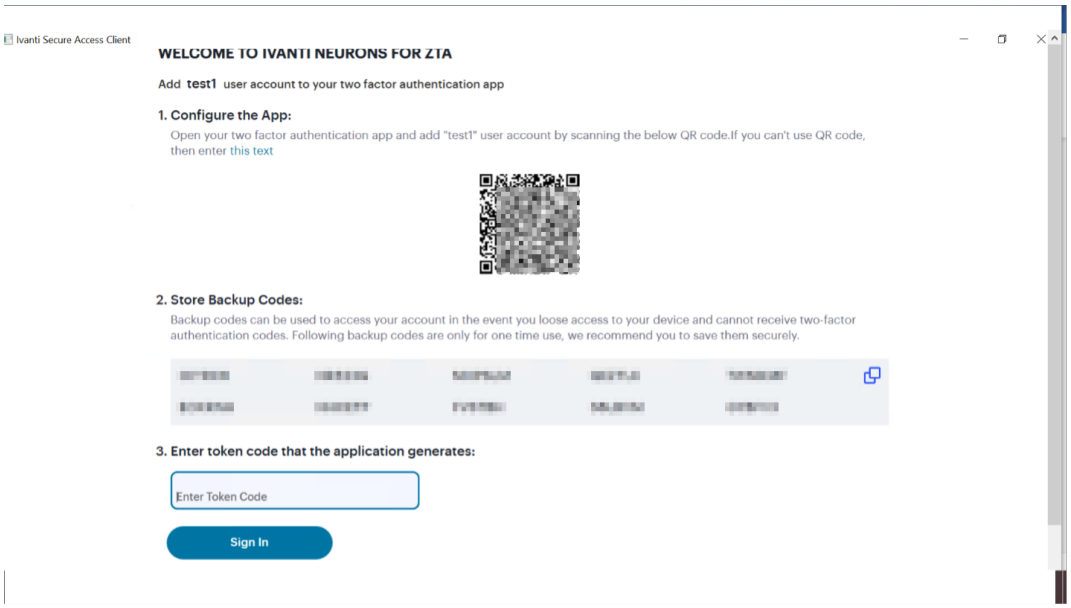

    First-time login TOTP registration

    Use this page to add your user details to an authenticator app on your device, such as Google Authenticator or Microsoft Authenticator. If you do not yet have such a two-factor authenticator app installed, do that now.

    Scan the QR code, or enter the provided text, to add the user details to your authenticator app. Then, store the generated backup codes in a secure location for future retrieval of a lost account.

    Finally, enter the token code generated by the authenticator app into the box provided, then select **Sign In**.

    For future sign-in attempts, the TOTP challenge dialog appears without the registration details or backup codes. As before, you provide the generated token code from your authenticator app into the box provided and select **Sign In**.

    When *Ivanti Secure Access Client* connects, it is minimised to the taskbar.

18. Open *Ivanti Secure Access Client* from the taskbar.

    Ivanti Secure Access Client appears. It shows the active connection to the *Controller* and presents a **Connect** button to access assigned applications and resources. For example:

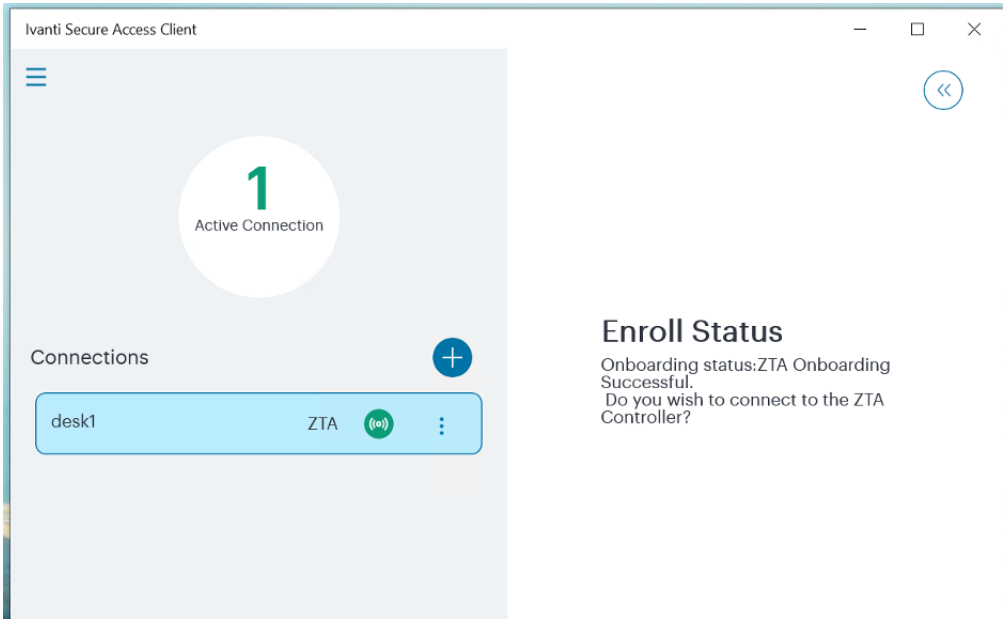

Ivanti Secure Access Client (Windows)

> To learn more about how *Ivanti Secure Access Client* maintains a connection with the *Controller*, and how user sessions are validated with your *nZTA Gateway*s, see Introduction.

19. In a typical enrollment, upon successful authentication to the *Controller*, *Ivanti Secure Access Client* automatically shows the end-user portal applications page through an embedded browser. For example:

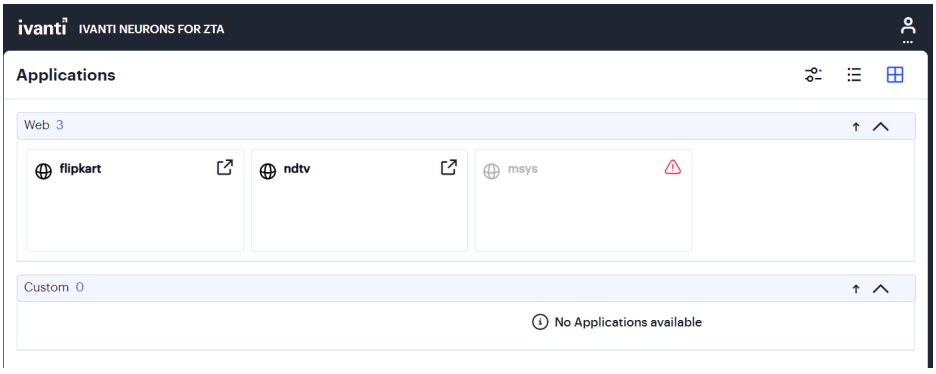

Assigned Applications and Resources (Windows)

After the *nZTA* end-user portal applications page appears, click any displayed resource to launch that item in your default system browser. To re-show the end-user portal at a future time, use the **ZTA** button in the *Ivanti Secure Access Client nZTA* connection.

When you launch an SSO (Single Sign-on) application from the end-user portal for the first time in a session, *nZTA* presents a pop-up dialog requesting the user to select a certificate with which to authenticate this device with the *Controller*. This is a one-time activity at the beginning of a session, and all further SSO application accesses (to any SSO application) re-use the same certificate.

If a default Gateway is configured on the *Controller*, and *nZTA* is the only active connection, the default Gateway handles all requests for unlisted applications from your Windows device. Refer to Using Application Discovery with Ivanti Secure Access Client.

Enabling Trusted Root CA Certificate on Windows Domain

To ensure that Windows desktop machines can successfully connect to the Windows domain, each machine must update its Trusted Root CA Certificates.

Ivanti recommends that the Domain Admin configures the Public Key Policies on the Group Policy Object and publishes it to all connected Windows desktops.

To add certificates to the Trusted Root Certification Authorities store for a Windows domain:

1. Access the Windows domain server and log in.

   Domain Admins is the minimum group membership required to complete this procedure.

2. Open the **Server Manager**.
3. Under **Features Summary**, click **Add Features**.
4. Select the **Group Policy Management** check box.
5. Click **Next**.
6. Click **Install**.

7. Wait until the **Installation Results** page shows that the installation of the **Group Policy Management Console** was successful.
8. Click **Close**.
9. Click **Start**, click **Administrative Tools** and then click **Group Policy Management**.

   The **Group Policy Management Console** appears.
10. In the console tree, double-click to expand *Group Policy Objects* in the forest and domain that contains the *Default Domain Policy* GPO.
11. Right-click the *Default Domain Policy* object, and then click **Edit**.
12. In the **Group Policy Management Console**, select *Computer Configuration* > *Windows Settings* > *Security Settings* > *Public Key Policies*.
13. Right-click the *Trusted Root Certification Authorities* store.
14. Click **Import** and follow the steps in the **Certificate Import Wizard** to import the certificates.

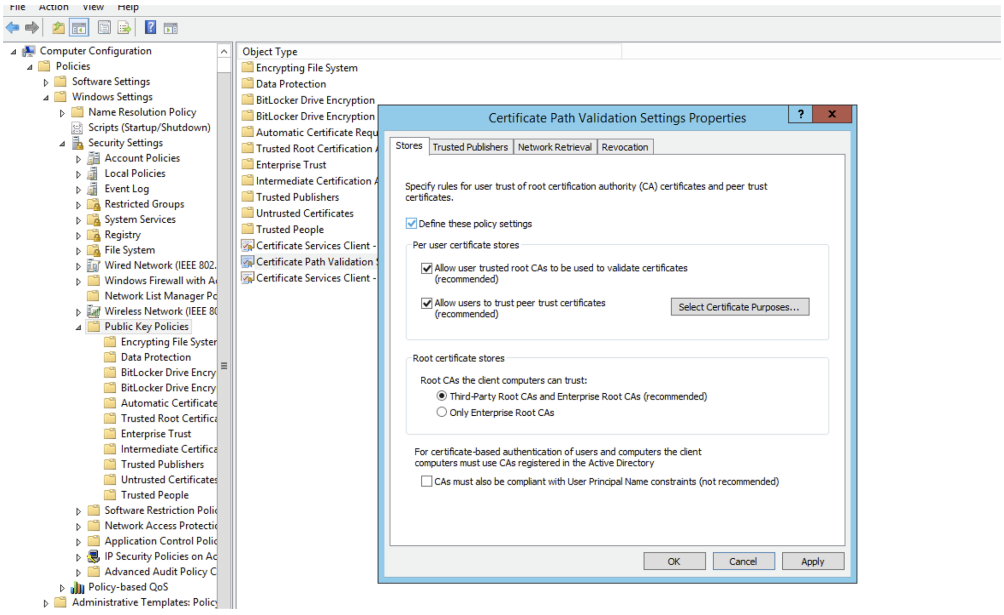

Enabling Trusted Root CA Certificate

When each Windows desktop machine next connects, the required Trusted Root CA Certificate installs automatically on the machine.

Enrolling a macOS Device

Before you start this process, you must have a sign-in URL for *nZTA*, based on the tenant FQDN provided by the *Ivanti* DevOps/Support organization.

> If you have an existing *Ivanti Secure Access Client* installed, you must first uninstall it before beginning the *nZTA* enrollment process.

To enroll a macOS desktop device:

1. Log into your macOS desktop.
2. Start a browser session.
3. In your browser, enter the *nZTA* sign-in URL into your address bar.

   For example,

   - Default: "https://tenant1.mycompany.com/login/enroll"
   - Custom: "https://tenant1.mycompany.com/login/<custom_user_signin>/".

   A login page appears.

4. Provide your credentials to access *nZTA* enrollment.

   The launcher page appears, for example:

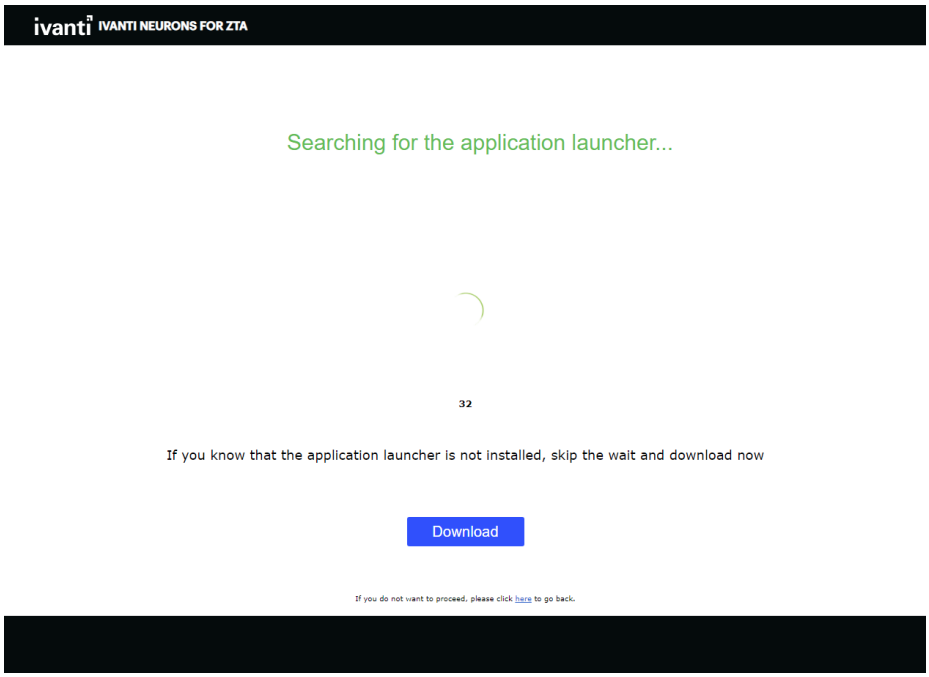

   nZTA Launcher macOS

5. Click **Download** and save the application launcher PulseSecureAppLauncher.dmg file locally.

   > A "HERE" hyperlink is displayed in the browser. This is required after the launcher has installed.

6. Click the downloaded file.

   The downloaded file opens, and a folder appears that contains the *PulseSecureAppLauncher.mpkg* file.

7. Double click the *PulseSecureAppLauncher.mpkg* file.

8. Click **Continue** when the install starts.

9. Wait until the install completes.

10. In the browser, click the "HERE" hyperlink to continue with the enrollment.

    A browser dialog requests confirmation to start the launcher app.

11. Click **Open PulseApplicationLauncher**.

    The launcher starts.

12. Confirm that you want to **Open** the application.

13. Confirm that you want the client to contact the *Controller*.

    Ivanti Secure Access Client then downloads and installs.

14. Enter your macOS device credentials.

    Ivanti Secure Access Client then installs and starts. The *nZTA* connection starts automatically. For example:

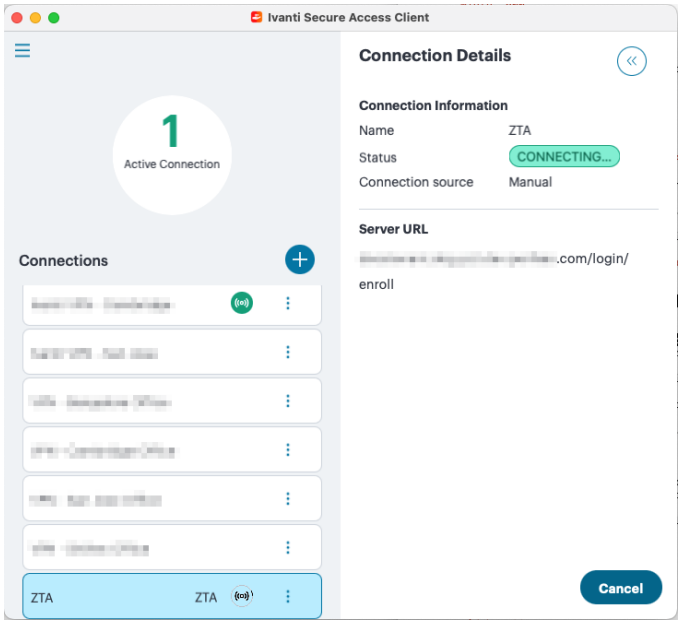

    macOS *Ivanti Secure Access Client*

> To learn more about how *Ivanti Secure Access Client* maintains a connection with the *Controller*, and how user sessions are validated with your *nZTA Gateways*, see Introduction.

    Ivanti Secure Access Client appears as an icon in the macOS system tray.

    Ivanti Secure Access Client in the System Tray (indicated)

    The connection activity reports on a number of tasks:

    - *Enrolling the User*.
    - *Fetching and Importing Client Certificates*. You must confirm any certificate requests.
    - *Fetching and Importing CA Certificates*. You must confirm any certificate requests.

- *Installing the CEF (Chromium Embedded Framework) browser*. This is an embedded browser used by *Ivanti Secure Access Client* for SAML-based login and to display the *nZTA* end-user portal applications page.

Onboarding is then complete.

15. In a typical enrollment, upon successful authentication to the *Controller*, *Ivanti Secure Access Client* automatically signs in the user.

16. (Optional) If your sign-in authentication policy is configured for Multi-Factor Authentication, you might be required to complete a TOTP (Time-based One Time Password) secondary authentication step each time you sign in.

    For a first time login, the user is presented with a TOTP registration page:

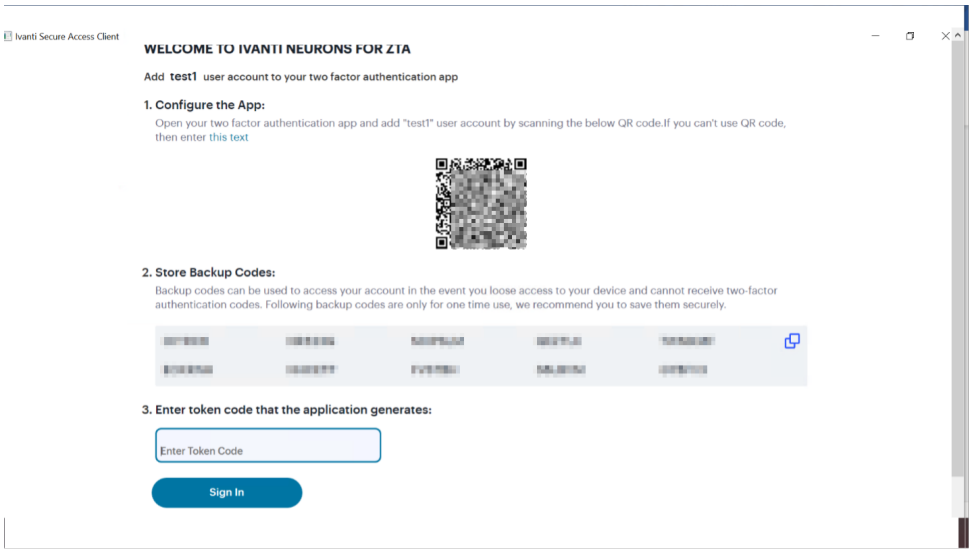

    First-time login TOTP registration

    Use this page to add your user details to an authenticator app on your device, such as Google Authenticator or Microsoft Authenticator. If you do not yet have such a two-factor authenticator app installed, do that now.

    Scan the QR code, or enter the provided text, to add the user details to your authenticator app. Then, store the generated backup codes in a secure location for future retrieval of a lost account.

    Finally, enter the token code generated by the authenticator app into the box provided, then select **Sign In**.

    For future sign-in attempts, the TOTP challenge dialog appears without the registration details or backup codes. As before, you provide the generated token code from your authenticator app into the box provided and select **Sign In**.

    When *Ivanti Secure Access Client* connects, it is minimised to the taskbar.

17. Accept the request to access the private key.

18. Accept the request to present a certificate for access to browser-based resources.

    The *nZTA* end-user portal applications page appears. For example:

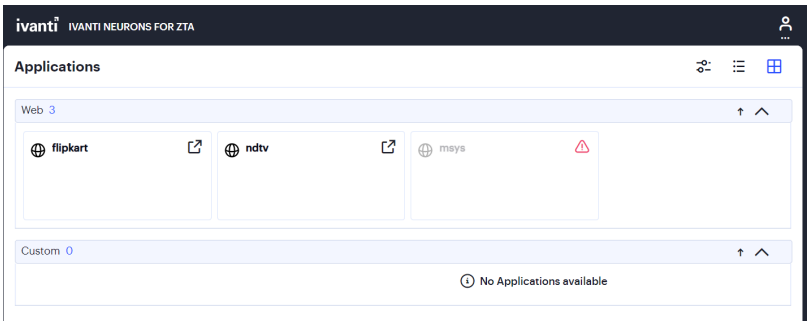

Assigned Applications and Resources

After the *nZTA* end-user portal applications page appears, click any displayed resource to launch that item in your default system browser. To re-show the end-user portal at a future time, click the **ZTA** button in the *Ivanti Secure Access Client nZTA* connection.

When you launch an SSO (Single Sign-on) application from the end-user portal for the first time in a session, *nZTA* presents a pop-up dialog requesting the user to select a certificate with which to authenticate this device with the *Controller*. This is a one-time activity at the beginning of a session, and all further SSO application accesses (to any SSO application) re-use the same certificate.

If a default Gateway is configured on the *Controller*, and *nZTA* is the only active connection, the default Gateway will handle all requests for unlisted applications from the macOS desktop device. Refer to Using Application Discovery with Ivanti Secure Access Client.

Enrolling a Linux Device

Browser-based enrollment is not supported for *Ivanti Secure Access Client* on Linux devices. Follow the instructions in this section to enroll the device by creating a *nZTA* connection through the *Ivanti Secure Access Client* application. To see the full feature support list for Linux devices, see Introduction.

Before you start this process, you must have:

- A Linux sign-in URL for *nZTA*, based on the tenant FQDN provided by the *Ivanti* DevOps/Support organization.
- The download location URL for your required installation package, as provided by the *Ivanti* DevOps/Support organization.

If you have an existing *Ivanti Secure Access Client* installed, you must first uninstall it before beginning the *nZTA* enrollment process.

Ivanti Secure Access Client is fully supported for use with *nZTA* on the following Linux variants:

- Ubuntu, see Enrolling on Ubuntu or Debian.

While not fully supported, the following Linux variants are considered compatible:

- Debian, see Enrolling on Ubuntu or Debian.
- Fedora, see Enrolling on Fedora or CentOS/RHEL.
- Centos/RHEL, see Enrolling on Fedora or CentOS/RHEL.

To learn more, contact your support representative.

Enrolling on Ubuntu or Debian

This section describes the installation of a *Ivanti Secure Access Client* Linux variants on either Ubuntu or Debian. If you want to install on either Fedora or CentOS/RHEL operating systems, refer to Enrolling on Fedora or CentOS/RHEL.

To enroll a Linux desktop device on Ubuntu or Debian:

1. Log into your Linux device.
2. Obtain the required *Ivanti Secure Access Client* for Linux installation package and download it to your Linux device. To obtain the installation package, contact your support representative.
3. Start a command line interface (CLI) terminal session.
4. If not already installed, install the prerequisite packages `nss3-tools` and `net-tools`. Make sure your package manager availability list is up to date, then run the following command for each package in turn:

   ```
   sudo apt-get install <dependency package name>
   ```

   Repeat this step for each prerequisite package.

5. Start the *Ivanti Secure Access Client* for Linux package installation at the prompt:

```
sudo dpkg -i <client installation package>
```

The installation process reports back to the session. For example, for Ubuntu:

```
Selecting previously unselected package pulsesecure.
(Reading database ... xxxxxx files and directories currently
installed.)
Preparing to unpack pulsesecure_9.1.R11_amd64.deb ...
Unpacking pulsesecure (9.1.R11) ...
Setting up pulsesecure (9.1.R11) ...
Created symlink /etc/systemd/system/multi-
user.target.wants/pulsesecure.service →
/lib/systemd/system/pulsesecure.service.
Processing triggers for desktop-file-utils (0.24-1ubuntu3)
...
Processing triggers for gnome-menus (3.36.0-1ubuntu1) ...
Processing triggers for mime-support (3.64ubuntu1) ...
Processing triggers for dbus (1.12.16-2ubuntu2.1) ...
Processing triggers for man-db (2.9.1-1) ...
```

6. Start the client. To do this, either start the *PulseUI* app from the **Activities** bar, or use the following from the command line:

```
/opt/pulsesecure/bin/pulseUI
```

The client appears:

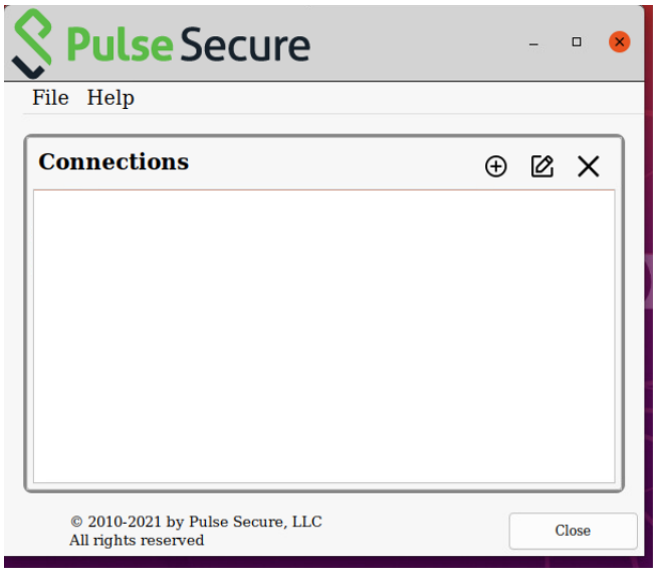

nZTA Ubuntu/Debian Linux Client

7. Add the required connection. To do this:

    - Click the *plus* icon in the client toolbar.

      The **Add Connection** dialog appears.

      

      nZTA Linux Add Connection

    - For **Type**, select *nZTA Connection*.
    - For **Name**, enter a suitably identifying name for the connection.
    - For **Server URL**, enter the Linux sign-in URL provided by the *Ivanti* DevOps/Support organization.

      For example,

        - Default: "https://tenant1.mycompany.com/login/enroll"
        - Custom: "https://tenant1.mycompany.com/login/<custom_user_signin>/".

- Click **Add** to add the connection and close the dialog.

  The new connection is added to the list of connections.

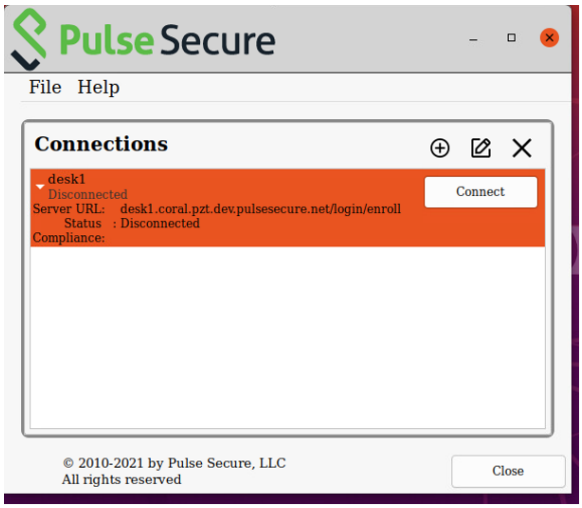

nZTA Linux nZTA Connection

8. For the *Ivanti Neurons for Zero Trust Access* connection, click **Connect**.

   A login dialog appears. For example:

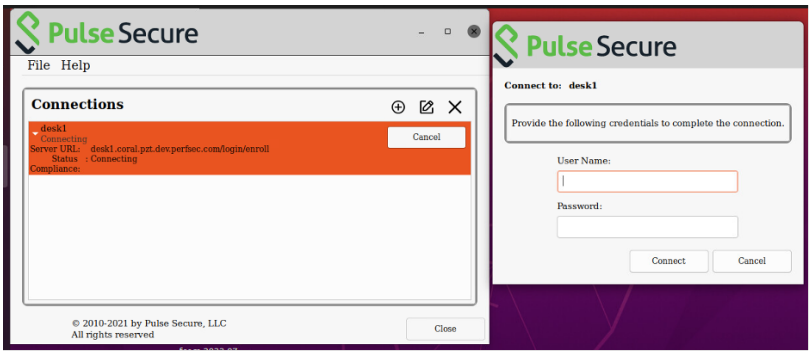

nZTA Linux nZTA Credentials

9. Enter your controller credentials and click **Connect**.

   A connection activity dialog reports a number of tasks:

   - *Enrolling the user*.
   - *Fetching and Importing of Client Certificates*. You must confirm any certificate requests.
   - *Fetching and Importing of CA Certificates*. You must confirm any certificate requests.
   - *Installing the CEF (Chromium Embedded Framework) browser*. This is an embedded browser used by *Ivanti Secure Access Client* for SAML-based login and to display the *nZTA* end-user portal applications page.
   - *Completion of the connection*. For example:

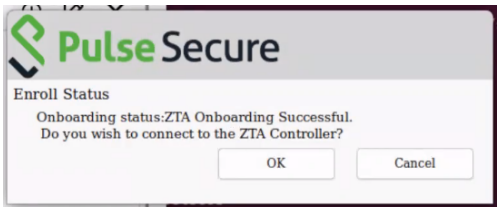

nZTA Linux ZTA Connection Success

10. Click **OK** to connect to *nZTA*.

    When this completes, the connection updates. For example:

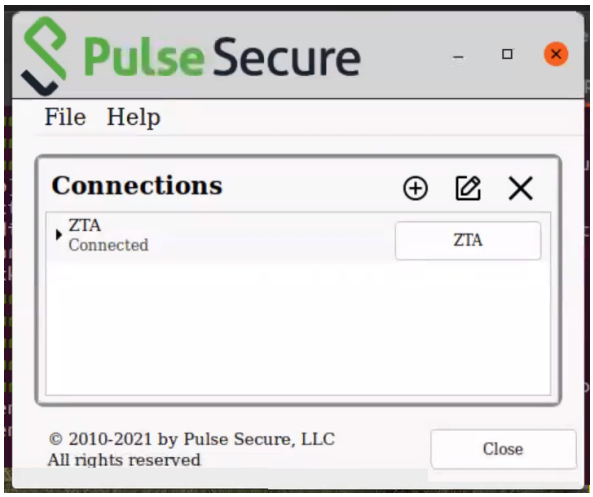

nZTA Linux ZTA Connected

Onboarding is then complete.

> To learn more about how *Ivanti Secure Access Client* maintains a connection with the *Controller*, and how user sessions are validated with your *nZTA Gateway*s, see Introduction.

11. (Optional) If your sign-in authentication policy is configured for Multi-Factor Authentication, you might be required to complete a TOTP (Time-based One Time Password) secondary authentication step each time you sign in.

    For a first time login, the user is presented with a TOTP registration page:

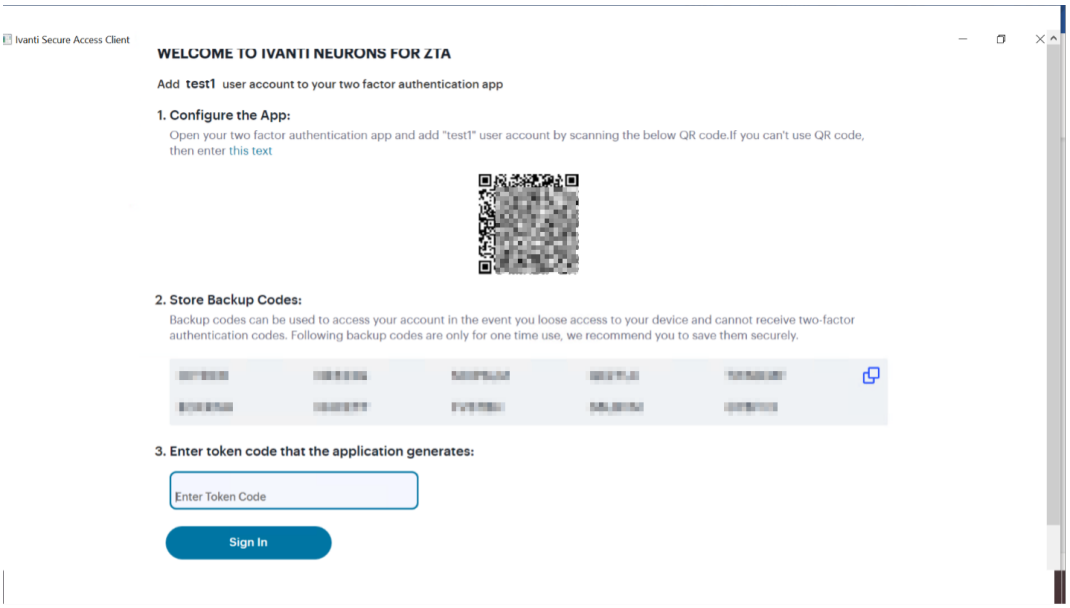

    First-time login TOTP registration

    Use this page to add your user details to an authenticator app on your device, such as Google Authenticator or Microsoft Authenticator. If you do not yet have such a two-factor authenticator app installed, do that now.

    Scan the QR code, or enter the provided text, to add the user details to your authenticator app. Then, store the generated backup codes in a secure location for future retrieval of a lost account.

    Finally, enter the token code generated by the authenticator app into the box provided, then select **Sign In**.

    For future sign-in attempts, the TOTP challenge dialog appears without the registration details or backup codes. As before, you provide the generated token code from your authenticator app into the box provided and select **Sign In**.

12. In a typical enrollment, upon successful authentication to the *Controller*, *Ivanti Secure Access Client* automatically shows the end-user portal applications page through an embedded browser. To re-show this portal at a future time, click the **ZTA** button in the *Ivanti Secure Access Client nZTA* connection. Alternatively, access your permitted applications from the Linux command line.

13. (Optional) To uninstall *Ivanti Secure Access Client* Linux variants, enter the following command at the prompt and provide a password:

```
[sudo] password for <user>: sudo apt-get purge pulsesecure
```

After the process completes successfully, the Linux client has been removed.

Ivanti Secure Access Client Linux variants do not support the use of default gateways.

Device rule types for *Ivanti Secure Access Client* Linux variants are limited to *File*, *Port*, and *Process*. For details of these device rule types, see Creating Device Policy Rules.

Enrolling on Fedora or CentOS/RHEL

This section describes the installation of *Ivanti Secure Access Client* Linux variants on either Fedora or CentOS/RHEL. If you want to install on either Ubuntu or Debian operating systems, refer to Enrolling on Ubuntu or Debian.

To enroll a Linux desktop device on Fedora or CentOS/RHEL:

1. Log into your Linux device.
2. Obtain the required *Ivanti Secure Access Client* for Linux installation package and download it to your Linux device. To obtain the installation package, contact your support representative.
3. Start a command line interface (CLI) terminal session.

4. If not already installed, install the prerequisite packages `nss3-tools` and `net-tools`. Make sure your package manager availability list is up to date, then run the following command for each package in turn:

   ```
   sudo yum install <dependency package name>
   ```

   Repeat this step for each prerequisite package.

5. Start the *Ivanti Secure Access Client* for Linux package installation at the prompt:

   ```
   sudo rpm -ivh <client installation package>
   ```

   The installation process reports back to the session. For example, for Fedora:

   ```
   Verifying...
   ################################# [100%]
   Preparing...
   ################################# [100%]
   Updating / installing...
   1:pulsesecure-2:9.1-R11
   ################################# [100%]
   rpm post../
   Created symlink /etc/systemd/system/multi-
   user.target.wants/pulsesecure.service →
   /lib/systemd/system/pulsesecure.service.
   Created symlink /etc/systemd/system/pulsesecure.service →
   /lib/systemd/system/pulsesecure.service.
   ```

6. Start the client. To do this, either start the *PulseUI* app from the **Activities** bar, or use the following from the command line:

```
/opt/pulsesecure/bin/pulseUI
```

The client appears:

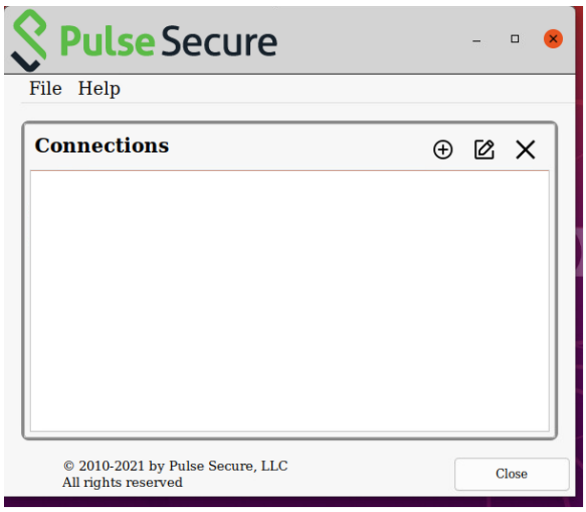

nZTA Ubuntu/Debian Linux Client

7. Add the required connection. To do this:

    - Click the *plus* icon in the client toolbar.

      The **Add Connection** dialog appears.

      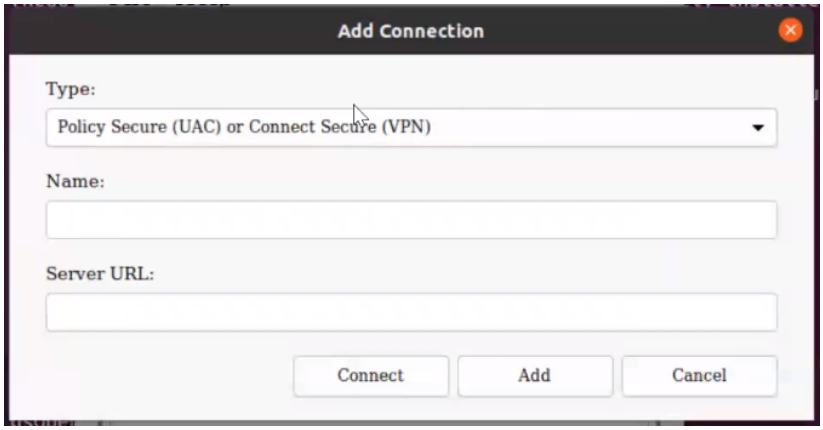

      nZTA Linux Add Connection

    - For **Type**, select *nZTA Connection*.
    - For **Name**, enter a suitably-identifying name.
    - For **Server URL**, enter the Linux sign-in URL provided by the *Ivanti* DevOps/Support organization.

      For example,

        - Default: “https://tenant1.mycompany.com/login/enroll”
        - Custom: “https://tenant1.mycompany.com/login/<custom_user_signin>/”

- Click **Add** to add the connection and close the dialog.

  The new connection is added to the list of connections.

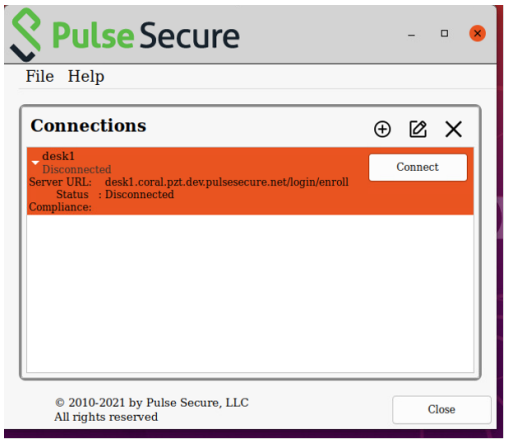

nZTA Linux nZTA Connection

8. For the *Ivanti Neurons for Zero Trust Access* connection, click **Connect**.

   A login dialog appears. For example:

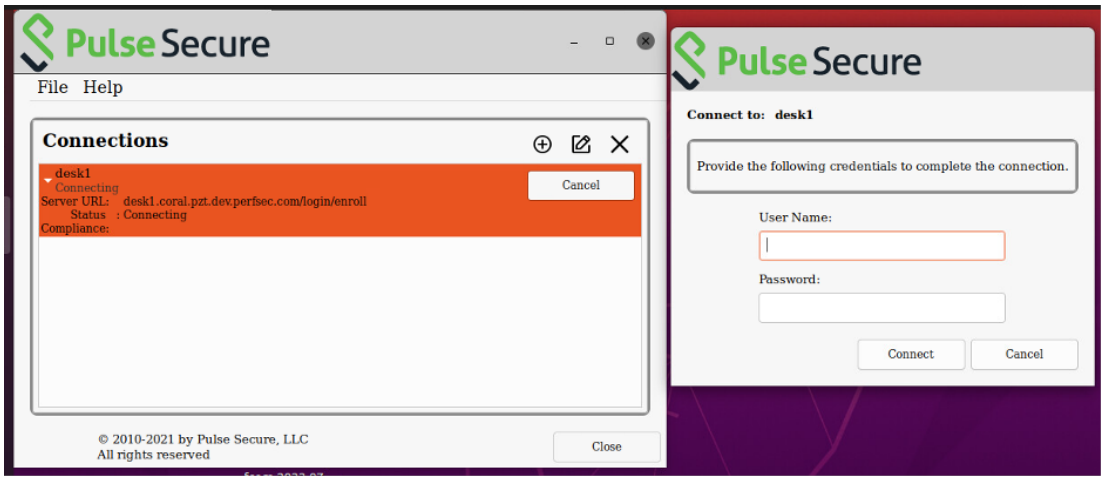

nZTA Linux nZTA Credentials

9. Enter your controller credentials and click **Connect**.

   A connection activity dialog reports a number of tasks:

   - *Enrolling the user*.
   - *Fetching and Importing of Client Certificates*. You must confirm any certificate requests.
   - *Fetching and Importing of CA Certificates*. You must confirm any certificate requests.
   - *Installing the CEF (Chromium Embedded Framework) browser*. This is an embedded browser used by *Ivanti Secure Access Client* for SAML-based login and to display the *nZTA* end-user portal applications page.
   - *Completion of the connection*. For example:

     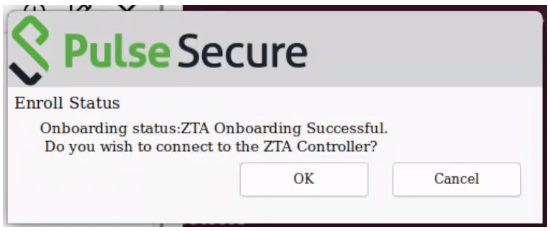

     nZTA Linux ZTA Connection Success

   Onboarding is then complete.

> To learn more about how *Ivanti Secure Access Client* maintains a connection with the *Controller*, and how user sessions are validated with your *nZTA Gateway*s, see Introduction.

10. (Optional) If your sign-in authentication policy is configured for Multi-Factor Authentication, you might be required to complete a TOTP (Time-based One Time Password) secondary authentication step each time you sign in.

    For a first time login, the user is presented with a TOTP registration page:

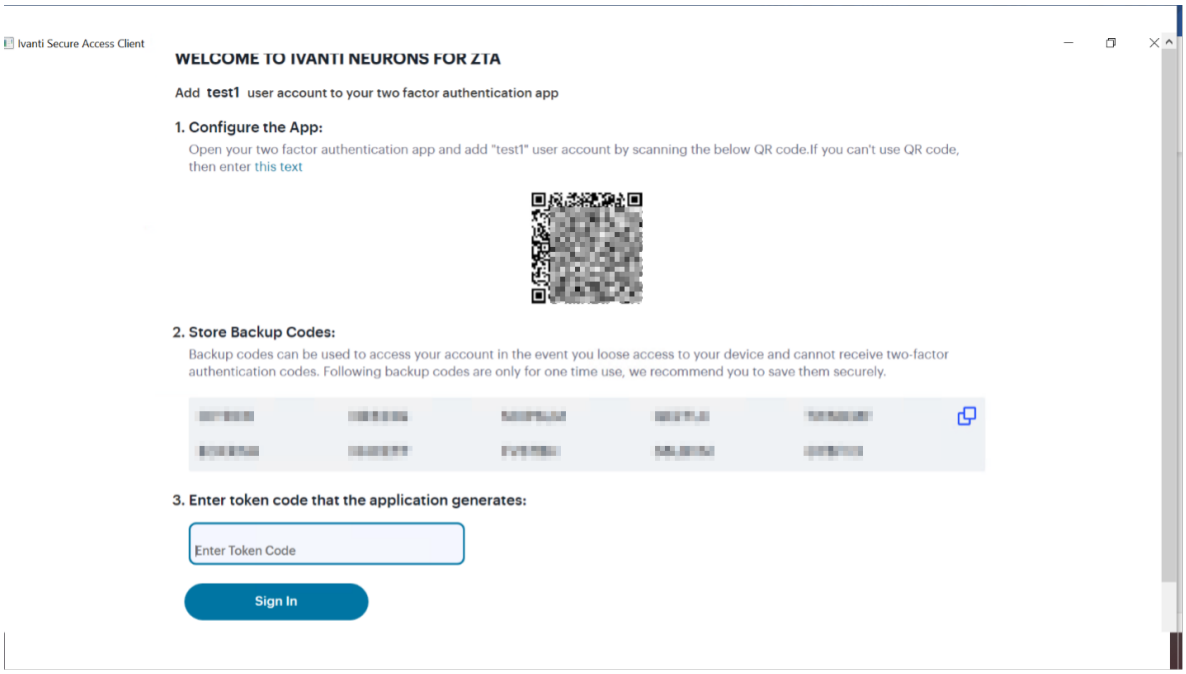

    First-time login TOTP registration

    Use this page to add your user details to an authenticator app on your device, such as Google Authenticator or Microsoft Authenticator. If you do not yet have such a two-factor authenticator app installed, do that now.

    Scan the QR code, or enter the provided text, to add the user details to your authenticator app. Then, store the generated backup codes in a secure location for future retrieval of a lost account.

    Finally, enter the token code generated by the authenticator app into the box provided, then select **Sign In**.

    For future sign-in attempts, the TOTP challenge dialog appears without the registration details or backup codes. As before, you provide the generated token code from your authenticator app into the box provided and select **Sign In**.

11. In a typical enrollment, upon successful authentication to the *Controller*, *Ivanti Secure Access Client* automatically shows the end-user portal applications page through an embedded browser. To re-show this portal at a future time, click the **ZTA** button in the *Ivanti Secure Access ClientnZTA* connection. Alternatively, access your permitted applications from the Linux command line.

12. (Optional) To uninstall *Ivanti Secure Access Client* Linux variants, enter the following command at the prompt and provide a password:

```
 sudo rpm -e pulsesecure
[sudo] password for <user>:
```

After the process completes successfully, the Linux client has been removed.

Ivanti Secure Access Client Linux variants do not support the use of default gateways.

Device rule types for *Ivanti Secure Access Client* Linux variants are limited to *File*, *Port*, and *Process*. For details of these device rule types, see Creating Device Policy Rules.

Enrolling an iOS Device

Application discovery and the use of a default Gateway is not supported from iOS mobile devices.

For mobile devices, *Ivanti Neurons for Zero Trust Access* compliance requires iOS v12.0 or later.

Before you start this process, you must have an iOS sign-in URL for *nZTA*, based on the tenant FQDN provided by the *Ivanti* DevOps/Support organization. The procedure to enroll your iOS device differs depending on whether you have an existing *Ivanti Secure Access Client* app installed and configured with a connection to a classic VPN product.

To enroll an iOS device that has a previous connection to a classic VPN product:

1. Start your iOS device and access its home page.

2. Locate and start the *Ivanti Secure Access Client* app.

3. In the main app menu, select **Enroll To ZTA**:

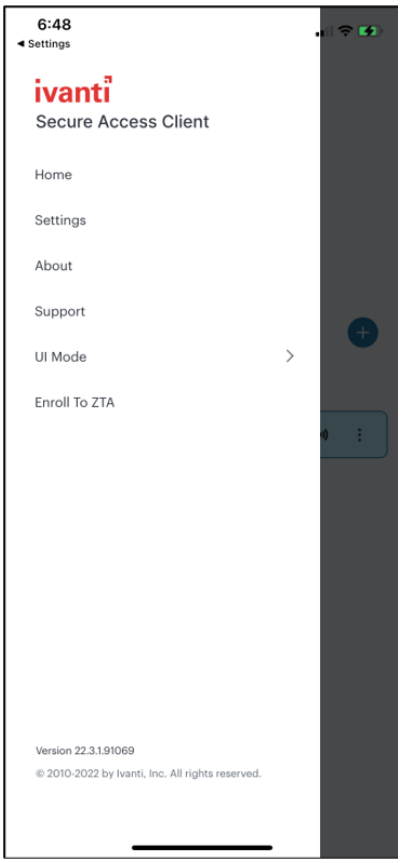

The *Ivanti Secure Access Client* app menu

The *Enroll with ZTA* screen appears.

4. The network type auto populates as Zero Trust Access. For Connection Name, specify a descriptive name for this connection. The name you specify appears in the Ivanti Secure Access Client interface.

5. For URL, specify the network that you want to connect to. Enter the nZTA controller URL as provided by the administrator.

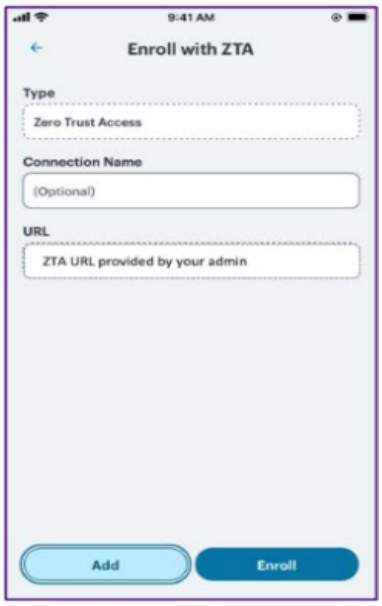

Enter sign-in URL

6. Click **Add** to save your new connection and the connection displays in the Home page. Click **Enroll** to add the connection and initiate a connection to the network.

7. To complete the enrollment procedure, follow the remaining steps below

To enroll an iOS device with no previous *Ivanti Secure Access Client* installation:

1. Start your iOS device and access its home page.

2. Open the App Store application.

3. In the search function, enter "Ivanti Secure Access Client".

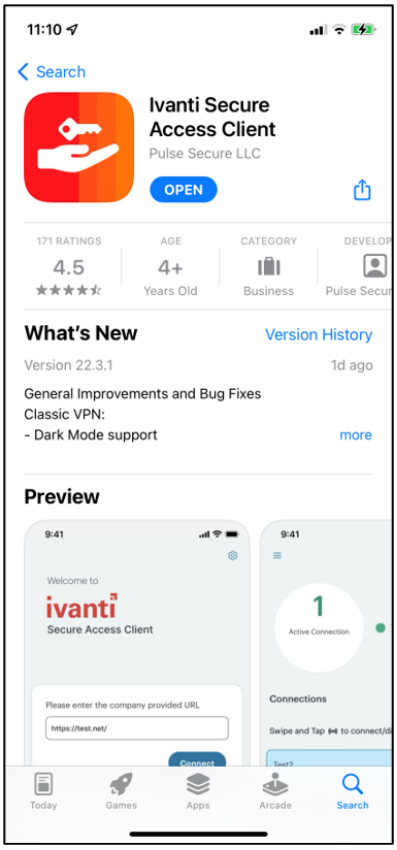

Searching the iOS App Store for *Ivanti Secure Access Client*

4. From the search results, locate, download, and install the *Ivanti Secure Access Client* app.

   The *Ivanti Secure Access Client* app is installed on your device and added to the interface. For example:

Ivanti Secure Access Client App Icon

5. Start the *Ivanti Secure Access Client* app.

   The welcome screen appears. For example:

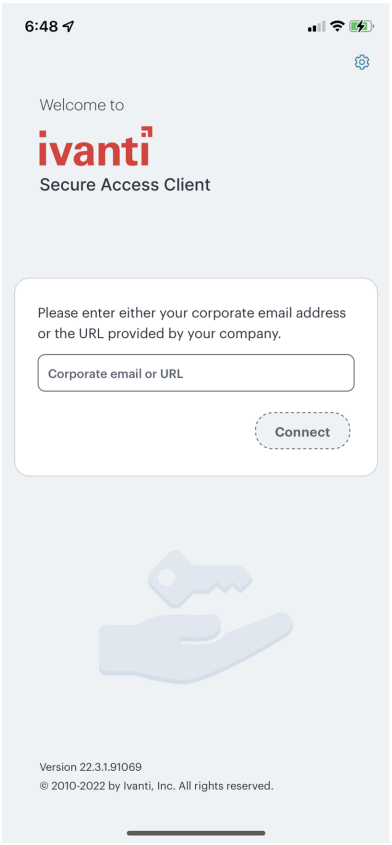

Ivanti Secure Access Client Welcome Screen

6. On the welcome screen, enter the *nZTA* sign-in URL provided in your invitation email.

   For example,

   - Default: “https://tenant1.mycompany.com/login/enroll”
   - Custom: “https://tenant1.mycompany.com/login/<custom_user_signin>/”.

7. Click **Connect**.
8. To complete the enrollment procedure, follow the remaining steps below

For both of the previous iOS device procedures, continue with the following steps:

1. Perform any required authentication for enrolling. (Local or Azure AD credentials)

   The **Connections** screen appears.

2. Accept any request to download a VPN configuration from the *Controller*. For example:

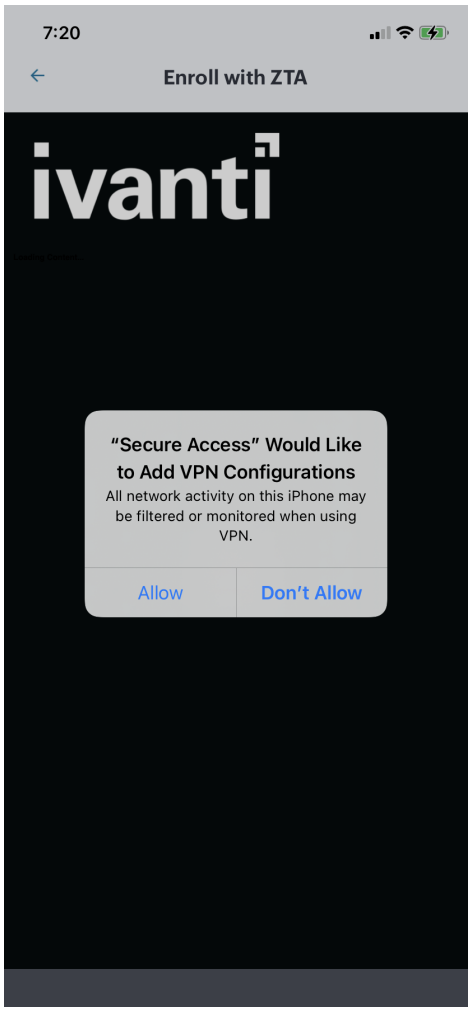

Adding VPN Configuration

After the *nZTA* profile is added to your Client configuration, the compliance of the device is checked. Then, session information is gathered and a VPN tunnel to the *Controller* is created automatically. This is indicated in the iOS status bar.

The **Connections** screen shows the active *Controller* connection:

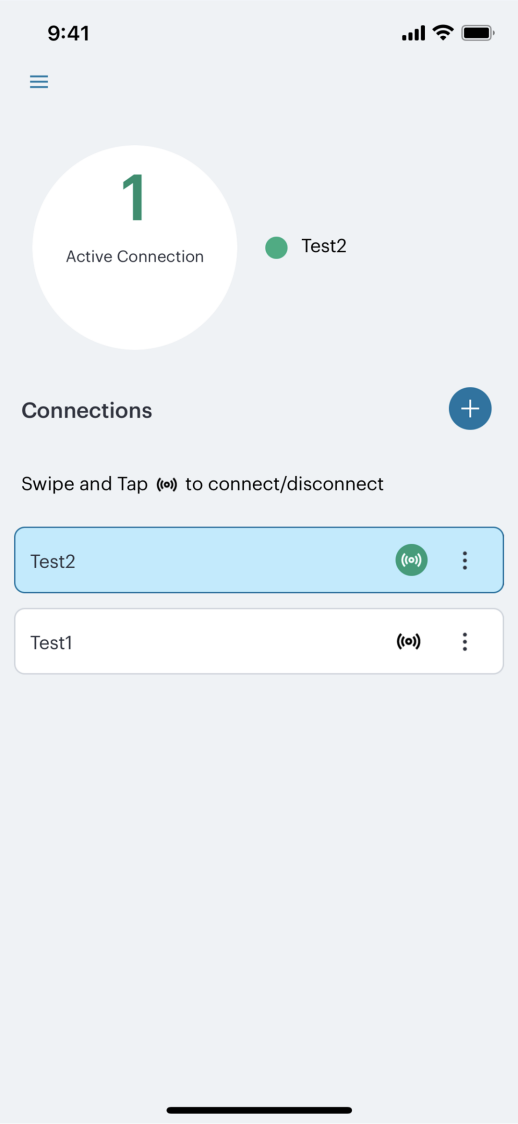

iOS Active *nZTA* Connection

To learn more about how *Ivanti Secure Access Client* maintains a connection with the *Controller*, and how user sessions are validated with your *nZTA Gateway*s, see Introduction.

3. (Optional) If your sign-in authentication policy is configured for Multi-Factor Authentication, you might be required to complete a TOTP (Time-based One Time Password) secondary authentication step each time you sign in.

   For a first time login, the user is presented with a TOTP registration page:

First-time sign in TOTP registration

Use this page to add your user details to an authenticator app on your device, such as Google Authenticator or Microsoft Authenticator. If you do not yet have such a two-factor authenticator app installed, do that now.

Scan the QR code, or enter the provided text, to add the user details to your authenticator app. Then, store the generated backup codes in a secure location for future retrieval of a lost account.

Finally, enter the token code generated by the authenticator app into the box provided, then tap **Sign In**.

For future sign-in attempts, the TOTP challenge dialog appears without the registration details or backup codes. As before, you provide the generated token code from your authenticator app into the box provided and select **Sign In**.

4. (Optional) Tap the **Disconnect** button to manually disable the *nZTA* connection. This facility overrides the on-demand connection feature and prevents *Ivanti Secure Access Client* from connecting to the *Controller* or any *nZTA Gateways*. By tapping this button, your *nZTA*-protected applications become inaccessible until the connection is restored. For more details, see Disabling the nZTA Connection.
5. (Optional) Tap the connection to view the connectivity status. For example:

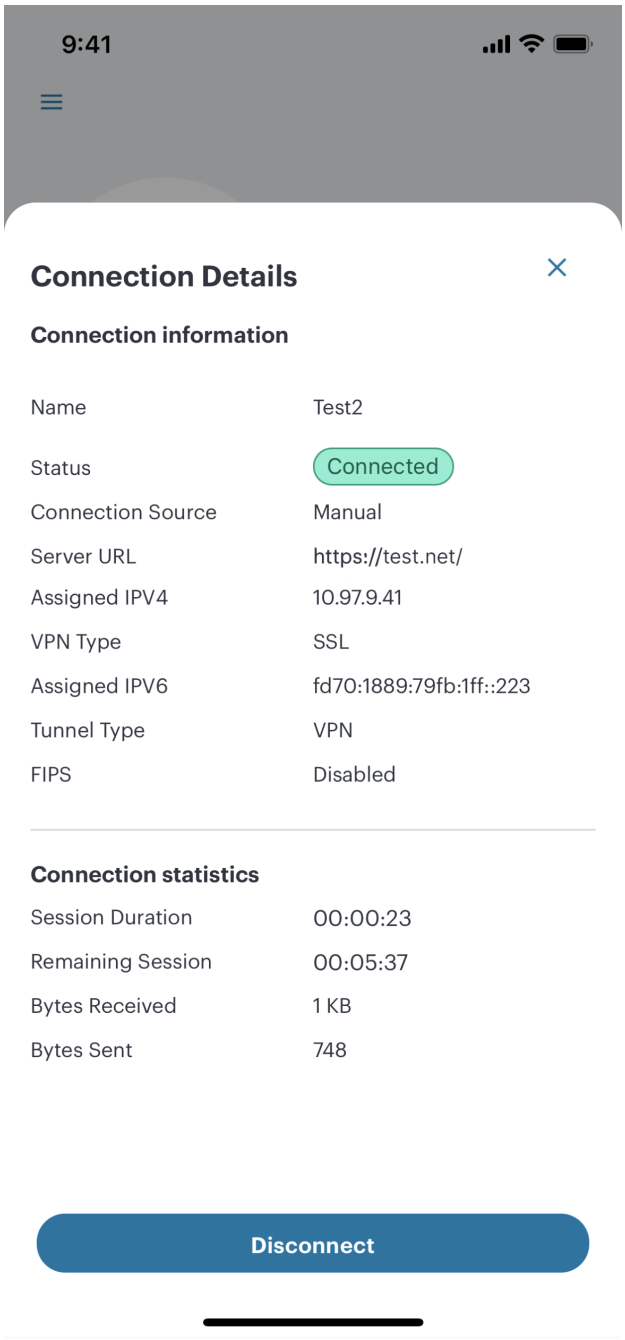

iOS Connectivity Status

To use a supported app, start the app as usual from the iOS interface, and enter any credentials if requested.

Installing a Beta Release of the iOS Client

Periodically, *Ivanti* might make available a pre-release *beta* version of *Ivanti Secure Access Client* for iOS for limited testing purposes. Beta releases of *Ivanti Secure Access Client* for iOS use the TestFlight package, a third-party app that enables users to download and test pre-GA packages of products. To learn more about TestFlight, see https://testflight.apple.com/.

Before you start, you must:

- Have an iOS sign-in/enrollment URL for *nZTA*, as provided by the *Ivanti* DevOps/Support organization.
- Install and register the TestFlight package from the Apple App Store.
- Configure TestFlight to have access to the *Ivanti Secure Access Client* app, using information provided by the *Ivanti* DevOps/Support organization.

To enroll a beta client on your iOS mobile device:

1. Locate and start the TestFlight app. For example:

   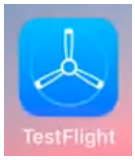

   TestFlight App Icon

2. In the TestFlight app, select the *Ivanti Secure Access Client* app and install it. For example:

   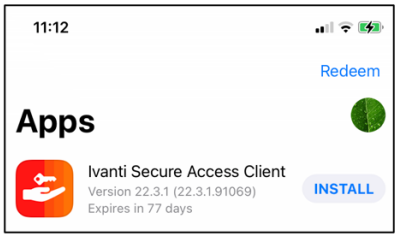

   TestFlight *Ivanti Secure Access Client* App Install

3. After the app installs, it is added to the iOS interface. For example:

   

   Ivanti Secure Access Client App Icon

4. Continue to enroll the device using the standard iOS Client enrollment procedure. Start the *Ivanti Secure Access Client* app and follow the remaining steps.

Enrolling an Android Device

Application discovery and the use of a default Gateway is not supported from Android devices.

For mobile devices, *Ivanti Neurons for Zero Trust Access* compliance requires Android v8.0 or later.

Before you start this process, you must:

- Have an Android sign-in URL for *nZTA*, based on the tenant FQDN provided by the *Ivanti* DevOps/Support organization.
- (For testing pre-GA/Beta *Ivanti Secure Access Client* packages only) Have an invitation email from your *Ivanti* representative that describes how to access pre-GA builds of the *Ivanti Secure Access Client* app from the Google Play store.

The procedure to enroll your Android device differs depending on whether you have an existing *Ivanti Secure Access Client* app installed and configured with a connection to a classic VPN product.

To enroll an Android device that has a previous connection to a classic VPN product:

1. Start your Android device and access its home page.
2. Locate and start the *Ivanti Secure Access Client* app.

3. In the main app menu, select **Enroll To ZTA**:

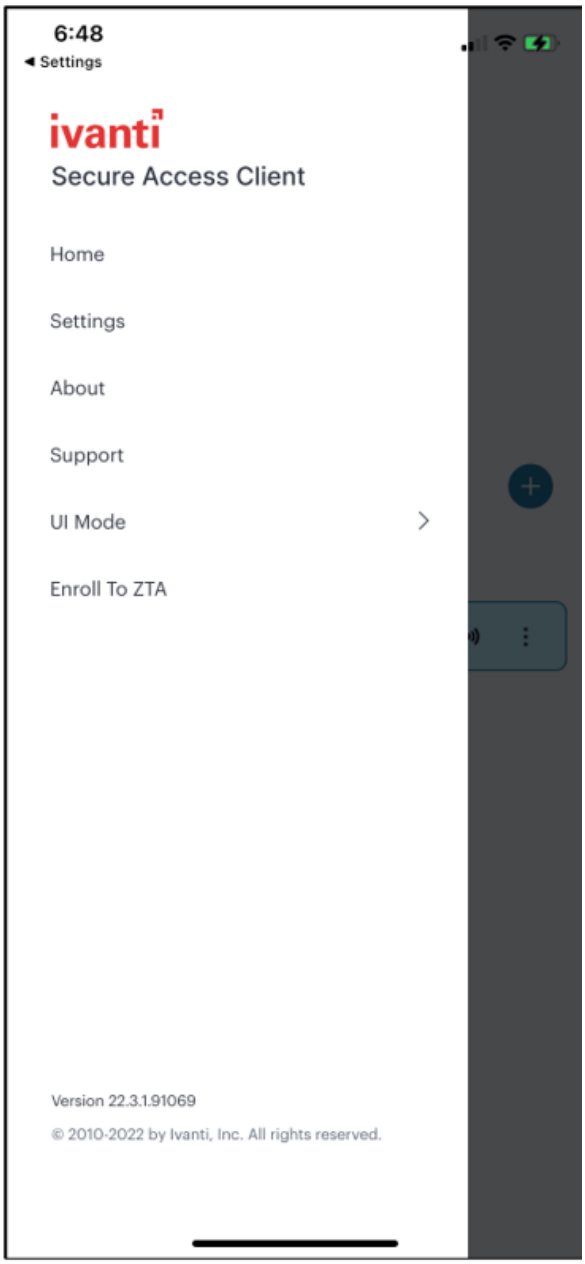

The *Ivanti Secure Access Client* app menu

The *Enroll with ZTA* screen appears.

4. The network type auto populates as Zero Trust Access. For Connection Name, specify a descriptive name for this connection. The name you specify appears in the Ivanti Secure Access Client interface.

5. For URL, specify the network that you want to connect to. Enter the nZTA controller URL as provided by the administrator.

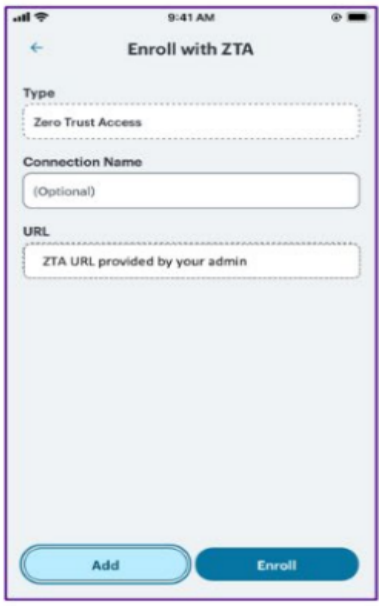

Enter sign-in URL

6. Click **Add** to save your new connection and the connection displays in the Home page. Click **Enroll** to add the connection and initiate a connection to the network.

7. To complete the enrollment procedure, follow the remaining steps below.

To enroll an Android device with no previous *Ivanti Secure Access Client* installation:

1. Start your Android device and access its home page.

2. Locate the *Ivanti Secure Access Client* app in the Google Play store and install it.

> If you are testing pre-GA/Beta builds of *Ivanti Secure Access Client* for Android, use the instructions provided in your invitation email to locate and download applicable pre-GA *Ivanti Secure Access Client* app builds in the Google Play store.

After the app installs, it is added to the Android interface. For example:

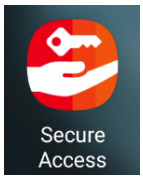

Ivanti Secure Access Client App Icon

3. Start the *Ivanti Secure Access Client* app.

   The welcome screen appears. For example:

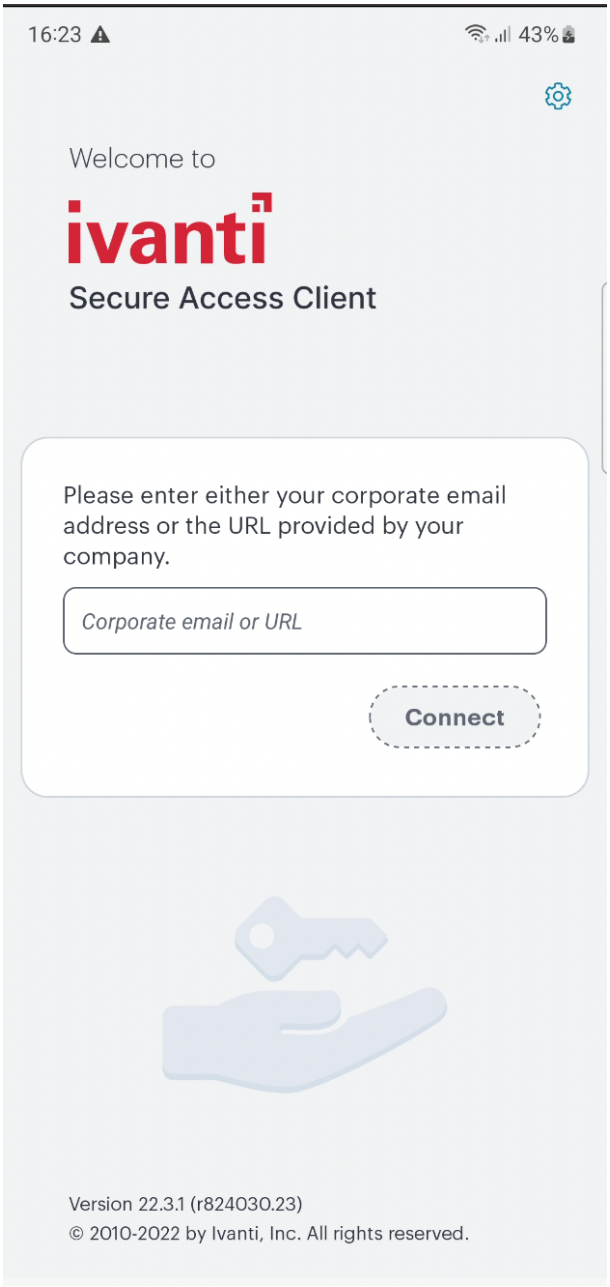

Ivanti Secure Access Client Welcome Screen

4. Enter the corporate email or *nZTA* sign-in URL.

   For example,

   - Default: “https://tenant1.mycompany.com/login/enroll”
   - Custom: “https://tenant1.mycompany.com/login/<custom_user_signin>/”.

5. Click **Connect**.
6. To complete the enrollment procedure, follow the remaining steps below.

_For both of the previous Android device procedures, continue with the following steps:

1. Perform any required authentication for enrolling. (Local or Azure AD credentials)

When using Azure AD authentication during onboarding, if you select "No" at the "Stay Signed In" prompt, you must re-enter your AD credentials whenever the login page prompts. To avoid this, make sure you select *Yes* at the **Stay Signed In** page.

The compliance of the device is checked, session information is gathered, and a VPN tunnel to the *Controller* is created automatically. This is indicated by a key icon and a *Ivanti Secure Access Client* icon.

The **Connections** screen shows the active *Controller* connection:

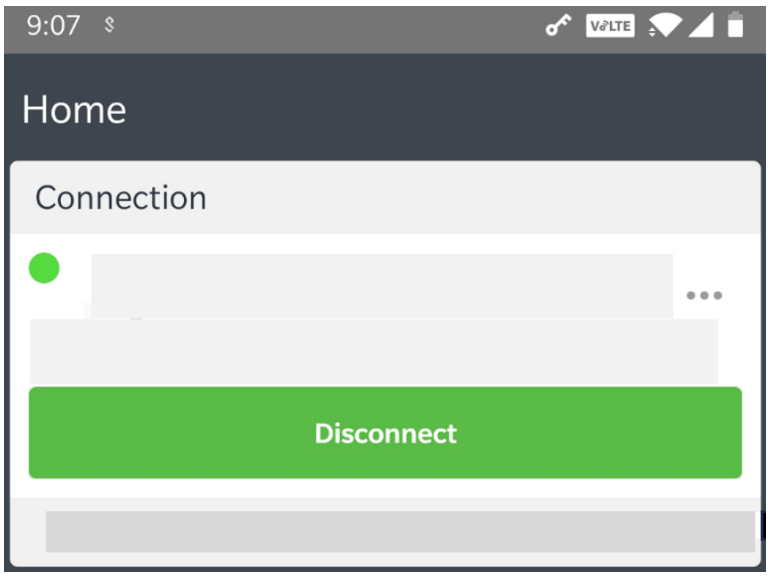

Android Active *nZTA* Connection

To learn more about how *Ivanti Secure Access Client* maintains a connection with the *Controller*, and how user sessions are validated with your *nZTA Gateways*, see Introduction.

2. (Optional) If your sign-in authentication policy is configured for Multi-Factor Authentication, you might be required to complete a TOTP (Time-based One Time Password) secondary authentication step each time you sign in.

   For a first time login, the user is presented with a TOTP registration page:

First-time sign in TOTP registration

Use this page to add your user details to an authenticator app on your device, such as Google Authenticator or Microsoft Authenticator. If you do not yet have such a two-factor authenticator app installed, do that now.

Scan the QR code, or enter the provided text, to add the user details to your authenticator app. Then, store the generated backup codes in a secure location for future retrieval of a lost account.

Finally, enter the token code generated by the authenticator app into the box provided, then tap **Sign In**.

For future sign-in attempts, the TOTP challenge dialog appears without the registration details or backup codes. As before, you provide the generated token code from your authenticator app into the box provided and select **Sign In**.

3. (Optional) Tap the **Disconnect** button to manually disable the *nZTA* connection. This facility overrides the on-demand connection feature and prevents *Ivanti Secure Access Client* from connecting to the *Controller* or any *nZTA Gateways*. By tapping this button, your *nZTA*-protected applications become inaccessible until the connection is restored. For more details, see Disabling the nZTA Connection.

4. (Optional) Tap the connection to view the connectivity status. For example:

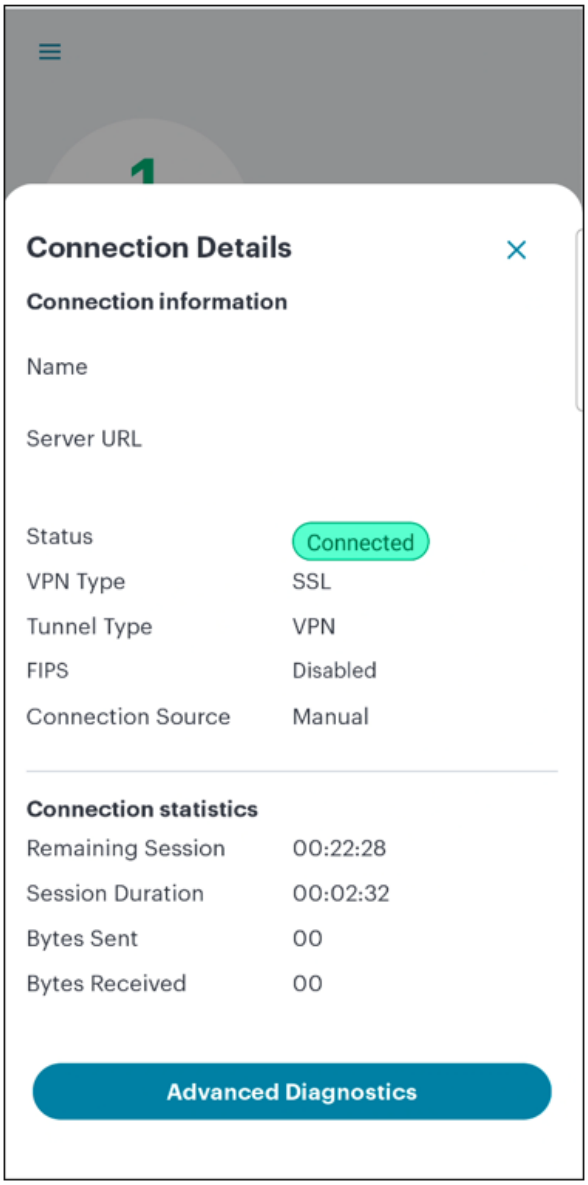

Android Connectivity Status

5. (Optional) Tap the page to view the Gateway Status. For example:

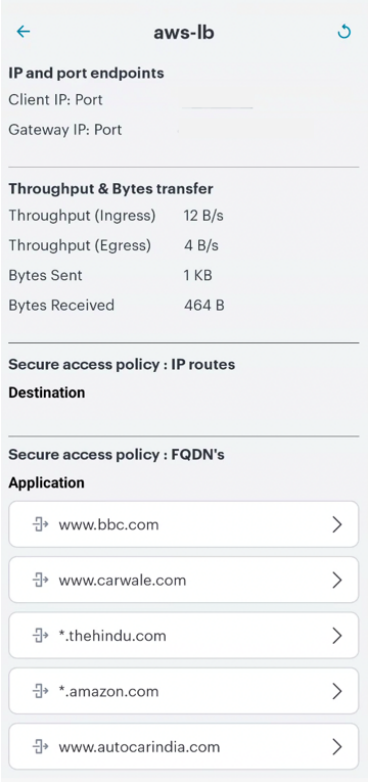

Android Gateway Status

6. (Optional) Tap the left arrow to return to the **Connections** screen.

To use a supported app, start the app as usual from the Android interface, and enter any credentials if requested.

Using *Ivanti Secure Access Client* with *nZTA*

Introduction

Ivanti provides a *nZTA*-ready version of the *Ivanti Secure Access Client* software required for end-user devices to be able to connect to your secure applications and resources.

Ivanti Secure Access Client connects to *nZTA* services, by default, through an on-demand connection basis and can handle multiple simultaneous *nZTA* and non-*nZTA* connections. To learn more, see On-Demand and Simultaneous Connection Handling.

Ivanti Secure Access Client maintains communication with the *Controller* to continuously-enable synchronization of policy and configuration updates. Through this mechanism, user requests to access resources and applications are subject to continuous assessment for risk and authorization. For more details, see Dynamic Policy Update and CARTA.

To learn more about enrolling user devices for use with *nZTA*, see Enrolling a User Device.

On-Demand and Simultaneous Connection Handling

While active, *Ivanti Secure Access Client* maintains two connection channels for *nZTA* services, a *control channel* to the *Controller*, and a *data channel* to your *nZTA Gateway*s. For more details on networking considerations when deploying *nZTA Gateway*s, see Working with Gateways.

The *control channel* connection to the *Controller* is activated when *Ivanti Secure Access Client* is started up and remains in an *always-on* state, silently in the background. If *Ivanti Secure Access Client* is able to locate a valid session cookie from an earlier session, the connection is re-established automatically. If no valid cookie is present, *Ivanti Secure Access Client* requests re-authentication from the user. The control channel is terminated when *Ivanti Secure Access Client* is shut down.

Ivanti Secure Access Client user sessions with a *nZTA Gateway* are subject to the following default timeout values:

- **5 minutes idle timeout**: If *Ivanti Secure Access Client* does not receive any traffic within this time period, the data channel to the connected Gateway is terminated, but the user session information is retained. If the user attempts the access the same resource(s) after this time (but before the maximum session timeout), the Gateway restores the data channel using the session information it currently holds.
- **60 minutes idle timeout**: If a *nZTA Gateway* receives no traffic from a connected client within this time period, the Gateway deletes the user session and instructs the *Controller* to do the same. If the user attempts to access the same resource(s) after this time, *Ivanti Secure Access Client* resumes the session with the Gateway automatically using cached credentials. However, if a Single Logout URL is specified in the SAML authentication methods covering user enrollment and sign-in, the user must re-authenticate manually for the connection to resume (see Using SAML Single Logout to Force User Authentication).

  

  User sessions for client connections with other Gateways are unaffected.

- **720 minutes session timeout**: If the user session reaches this maximum length, the session is terminated and re-authentication is requested from the client.

These values are the default settings. To configure custom *idle timeout* and *max session length* values for your user sessions, see Configuring Session Timeouts.

Ivanti Secure Access Client creates *data channel* connections to *nZTA Gateway*s as an *on-demand* service. That is, connections to resources and applications controlled by *nZTA Gateway*s become active only when required, and the connection is suspended after a period of inactivity. The user remains unaware of the connection state, unless re-authentication becomes necessary. As a user makes a request for a resource, *Ivanti Secure Access Client* transitions automatically from disconnected to connected. The connection remains in this state for the duration of the session, or until one of the following events occurs:

- An idle time-out occurs (after 5 minutes)
- The connection is actively placed in a disconnected state
- *Ivanti Secure Access Client* is shut down

To avoid the data channel being reconnected unnecessarily, non-*nZTA* DNS traffic is redirected to the device's physical network adapter.

Applicable *Ivanti Secure Access Client* versions can manage simultaneous connections with the *Controller*, and with other *Ivanti* services such as *Ivanti Connect Secure* (*ICS*). While *ICS* connections must be activated and deactivated by the user, connections to *nZTA* are provided on-demand, as mentioned. Therefore, a *nZTA* connection in the *Ivanti Secure Access Client* does not provide the same **Connect** and **Disconnect** controls. Instead, *nZTA* connections include only a **ZTA** button to provide access to the *nZTA* Applications page. If this button is active, the connection to the *Controller* has been established. If the button is inactive, the connection to the *Controller* has not yet been established, or a communication problem has occurred. In this case, access to your applications is prevented.

In a simultaneous dual connection scenario:

The standard VPN connection must be started first. Then, the *nZTA* VPN connection can be started.

- If the *nZTA* VPN connection is started first, a standard VPN connection *cannot* be started afterwards if a default Gateway is in use, see Application Discovery with Ivanti Secure Access Client.
- All requests from applications referenced in secure access policies are handled by *nZTA Gateway*s. All other requests will be processed over the standard VPN outside *nZTA* , and will potentially be insecure.
- When simultaneous dual VPN connections are active, application discovery using a default Gateway is *not* supported.

Simultaneous dual connections are not currently supported on Linux clients.

When using simultaneous dual connections, if the standard VPN is subsequently disabled, application discovery and a default Gateway will be activated automatically if configured. Currently, a default Gateway can only be utilized by requests from applications on macOS/Windows desktop devices.

When running active connections to both *nZTA* and *ICS* simultaneously, note that the following *ICS* features are not supported:

- Route Monitoring
- Traffic Enforcement
- Stealth Mode
- Always on VPN/LockDown
- Location awareness
- IPv6 support

On-demand connections are not currently supported on Linux clients.

Resource Precedence Over Simultaneous Connections

In some cases, the simultaneous connection handling feature in *Ivanti Secure Access Client* can mean that FQDN-based or IP address-based resources can be configured across multiple *nZTA* and non-*nZTA Gateways* (such as *ICS*) simultaneously. Where such conflicts occurs, *Ivanti Secure Access Client* evaluates the order of precedence before establishing the connection.

The following table describes the order of precedence from highest to lowest:

Order of Precedence for Resource Access

Priority	Order of Application Precedence	Remarks
1	*nZTA* FQDN resource	*nZTA* FQDN resources have the highest precedence
2	*nZTA* IP address resource - with Deny Policy	
3	*nZTA* IP address resource - with Allow Policy	
4	Non-*nZTA* FQDN/IP Address resources	Non-*nZTA* resources have the lowest precedence in simultaneous connection handling mode

The following table describes various resource conflict use-cases and the resolution applied by *nZTA* based on the above order of precedence:

Resource Conflict Use Cases

Use Case	Resource Conflict	Resolution
1	A *nZTA* FQDN resource conflicts with a *nZTA* IP Allow resource or *nZTA* IP Deny resource.	The *nZTA* FQDN resource is given precedence.
2	A *nZTA* IP Deny resource conflicts with a *nZTA* IP Allow resource.	The *nZTA* IP Deny resource is given precedence.
3	A *nZTA* IP Deny resource through one *nZTA*	The first *nZTA Gateway* IP resource takes

Use Case	Resource Conflict	Resolution
	Gateway with a *nZTA* IP Deny Resource in another *ZTA Gateway*.	effect as both are of same precedence.
4	A *nZTA* IP Allow resource through one *nZTA Gateway* with a *nZTA* IP Allow resource in another *ZTA Gateway*.	The first *nZTA Gateway* IP resource takes effect as both are of same precedence.
5	A *nZTA* FQDN resource conflicts with a non-*nZTA* FQDN or IP resource.	The *nZTA* FQDN resource is given precedence.
6	A *nZTA* IP Deny or Allow resource conflicts with a non-*nZTA* IP resource.	The *nZTA* IP resource is given precedence.
7	A *nZTA* single IP resource is conflicted to a *nZTA* subnet range of IP addresses.	Both routes are added to the client device routing table and the Operating System behavior is leveraged to choose the longest prefix match. Routing is established based on this criteria.
8	A *nZTA* single IP/subnet resource is conflicted to a non-*nZTA* subnet range of IP addresses	Both routes are added to the client device routing table and the Operating System behavior is leveraged to choose the longest prefix match. Routing is established based on this criteria.
9	A non-*nZTA* include FQDN resource is also configured as a *nZTA* FQDN resource.	**[Split tunnel]** First connect to the non-*nZTA* endpoint, and then to the *nZTA* endpoint (establishing simultaneous connection handling mode). Then, access the resource. The *nZTA* FQDN is given precedence.
10	A non-*nZTA* include FQDN Resource is also configured as a *nZTA* FQDN resource.	**[Split tunnel]** First connect to the *nZTA* endpoint, and then connect to the non-*nZTA* endpoint (establishing simultaneous connection handling mode). Then, access the resource. The *nZTA* FQDN is given precedence.

The following scenarios are exceptions and can result in unexpected behavior:

Exceptions

Case	Resource Conflict	Remarks
1	A *nZTA* FQDN or IP resource conflicts with a DNS IP address	Currently, precedence is not given to the DNS IP address configuration. This can result in unexpected behavior.
2	A *nZTA* IP resource conflicts with local subnet routes in a non-*nZTA* connection with local subnet route precedence.	The *nZTA* IP resource is given precedence and the non-*nZTA* endpoint does not operate as expected in simultaneous connection handling mode.
3	A *nZTA* IP resource conflicts with a non-*nZTA* server URI.	If the non-*nZTA* server URI is resolved to an IP address which is configured as a *nZTA* IP resource, traffic is directed through *nZTA* due to the higher precedence.

Using SAML Single Logout to Force User Authentication

By default, *Ivanti Secure Access Client* attempts to re-authenticate a connection to a *nZTA*-controlled resource using cached credentials for the session. In other words, during a single valid session, a user is not required to re-enter their credentials each time the on-demand connection is reactivated following an idle timeout or manual disconnection.

> To learn more about on-demand client connections, see On-Demand and Simultaneous Connection Handling.

An organization might want to establish a policy of increased security for connecting devices, such that prior to reconnecting the user is required to re-authenticate manually with the SAML IdP. To support this, *nZTA* includes the ability to specify a single logout/sign-out URL in the SAML authentication method you use for user enrollment or login.

By specifying a Single Logout URL in your SAML user authentication methods, your users must re-enter their credentials each time the *nZTA* connection is re-established.

To add a Single Logout URL to an existing user authentication method:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Secure Access** icon, then select **Manage Users > Authentication Servers**.

   The **Authentication Servers** page appears. This page lists all existing user authentication methods. For example:

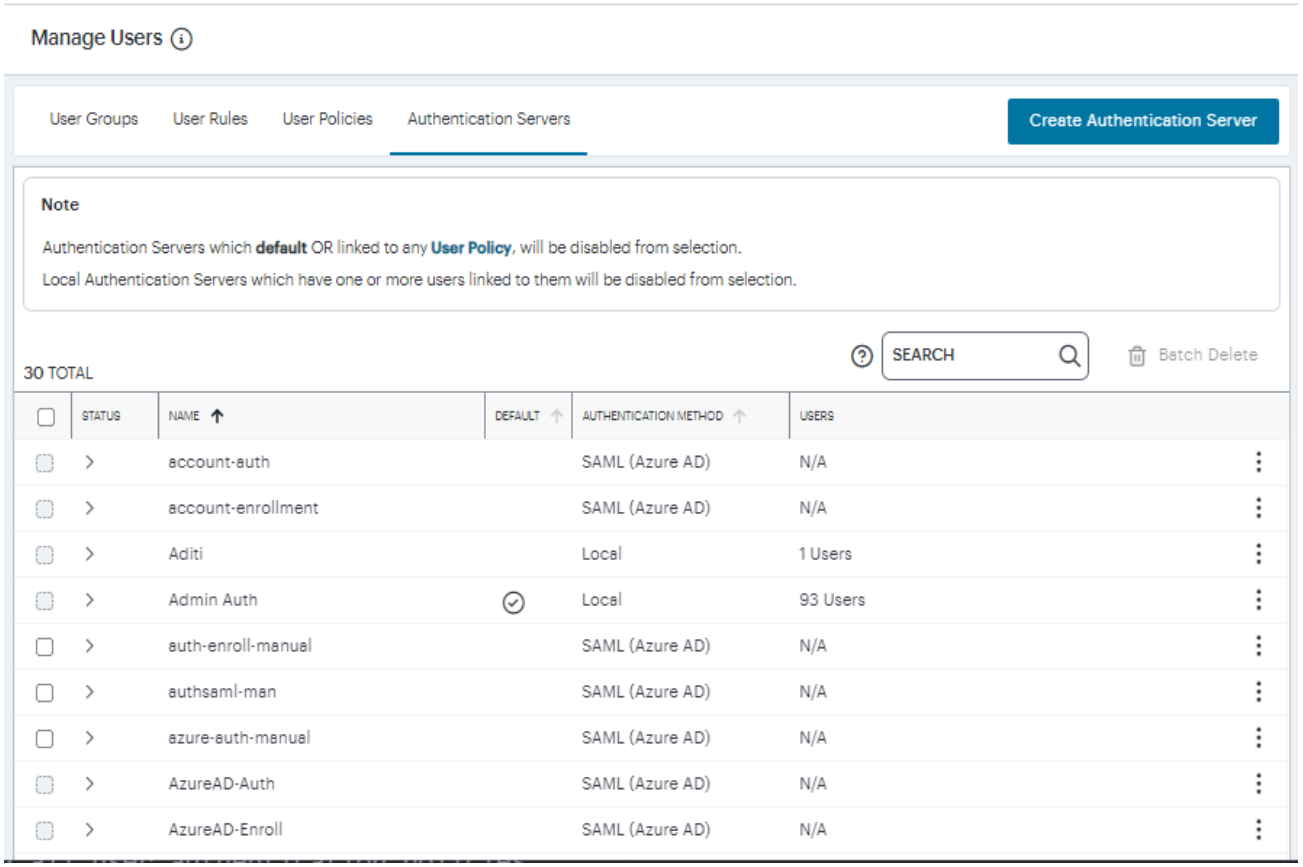

   User Authentication Methods

3. Click the three dots adjacent to the user authentication method you want to modify, and click **Edit**.

   The *Edit Authentication Method* dialog appears.

4. Specify your **Single Logout URL** in the text box provided:

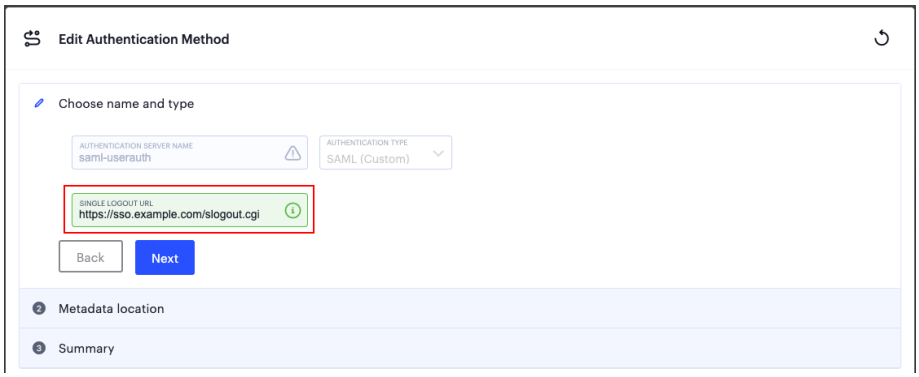

Adding a Single Logout URL to a SAML authentication method

1. To save your changes, click **Save Changes**.
2. Click **Next**, then **Next** again to view the *Summary* step.

Conversely, to re-establish the default service behavior of not requiring manual authentication upon re-connection, repeat these steps and remove the **Single Logout URL** from your existing authentication method.

To learn more about configuring an authentication method, see Working with User Authentication.

Disabling the *nZTA* Connection

Ivanti Secure Access Client additionally provides the ability to actively disable the on-demand connection feature. Use of this facility disables the *nZTA* connection, avoiding the scenario where *Ivanti Secure Access Client* attempts to repeatedly request authentication even after the user might be unable to authenticate due to too many failed attempts, or where the user just does not require access to any *nZTA*-controlled resources during that session.

If a user attempts to request a *nZTA*-controlled resource during the period a *nZTA* connection is disabled, the request fails. Other *Ivanti Secure Access Client* connections are unaffected.

To disable a *nZTA* connection, use one of the following mechanisms:

- For *Ivanti Secure Access Client* on macOS and Windows, click **Disconnect** in the *Ivanti Secure Access Client* connection list context menu. Right-click a *nZTA* connection profile to see the available options.

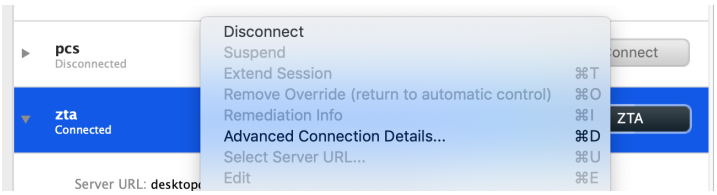

Manually disabling a *nZTA* connection through the *Ivanti Secure Access Client* connection list context menu

- For *Ivanti Secure Access Client* on macOS and Windows, click **Disconnect** from the *System Tray* client icon. View the sub-menu for the *nZTA* connection you want to disconnect.

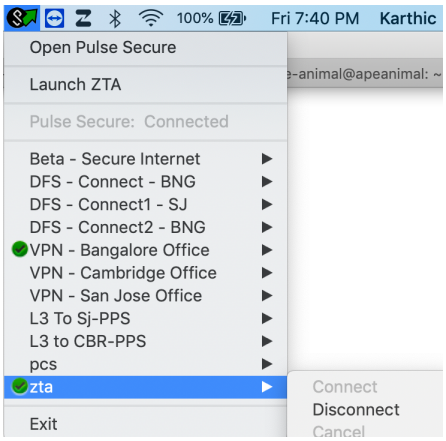

Manually disabling a *nZTA* connection through the *Ivanti Secure Access Client* System Tray control

- For *Ivanti Secure Access Client* on macOS and Windows, click **Disconnect** through the *Ivanti Secure Access Client* application menu. Open *Ivanti Secure Access Client* and select the *nZTA* connection profile. Then click **File > Connections > Disconnect**.

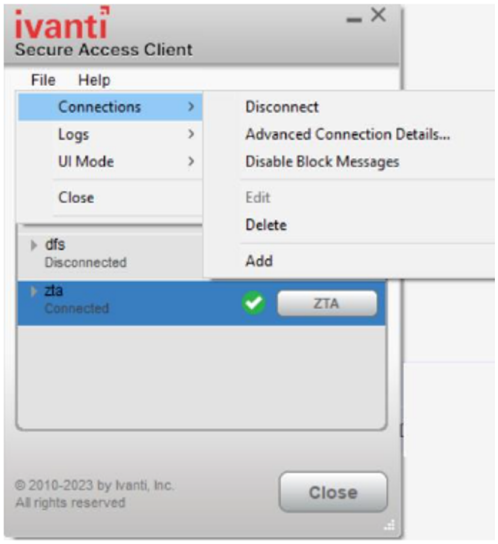

Manually disabling a *nZTA* connection through the *Ivanti Secure Access Client* application menu

- For *Ivanti Secure Access Client* on Android and iOS devices, use the **Disconnect** option in the *Ivanti Secure Access Client* application. Open *Ivanti Secure Access Client*, locate the *nZTA* connection profile, and tap the **Disconnect** button.

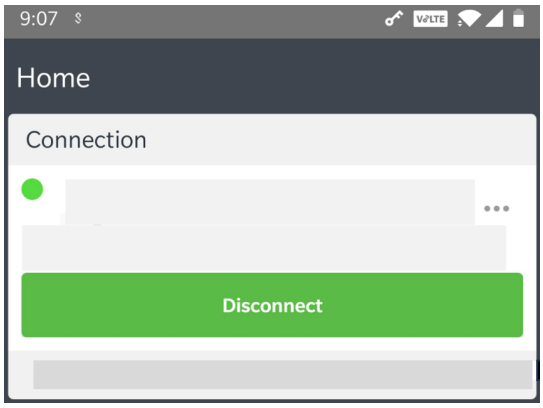

Manually disabling a *nZTA* connection through the *Ivanti Secure Access Client* application (Android device example)

By setting the *nZTA* connection to be disconnected, *Ivanti Secure Access Client* suspends both the *control channel* and the *data channel* (where either are active). If the *control channel* was previously logged-in to the *Controller*, this remains the case to facilitate session resumption through a subsequent reconnect.

The disconnect feature is not activated by clicking or tapping **Cancel** in the *nZTA* authentication dialog. Canceling an authentication request triggers a timeout interval, after which *Ivanti Secure Access Client* re-displays the authentication dialog. The disconnect feature instead disables the authentication request process until the user manually reinstates it.

To reinstate the *nZTA* connection on macOS and Windows devices, use the **Launch ZTA** option in the *Ivanti Secure Access Client* system tray menu:

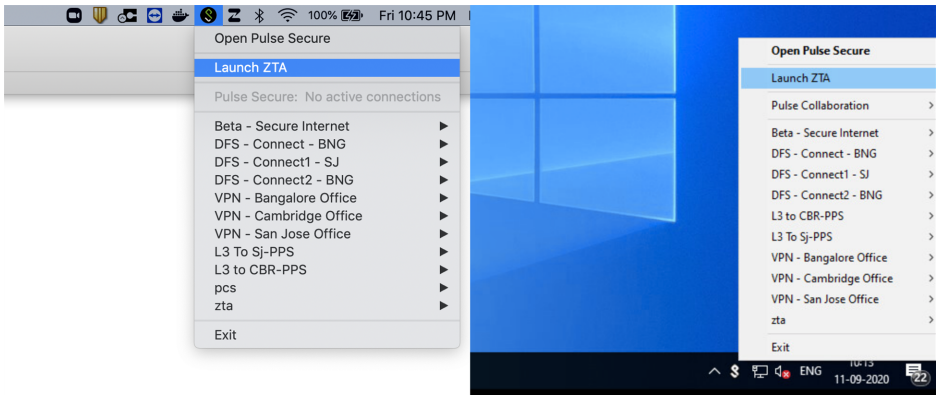

Restarting a *nZTA* connection through the *Ivanti Secure Access Client* System Tray menu (macOS and Windows variants)

To reinstate the *nZTA* connection on Android and iOS devices, tap the **ZTA** button in the *nZTA* connection profile in the *Ivanti Secure Access Client* application:

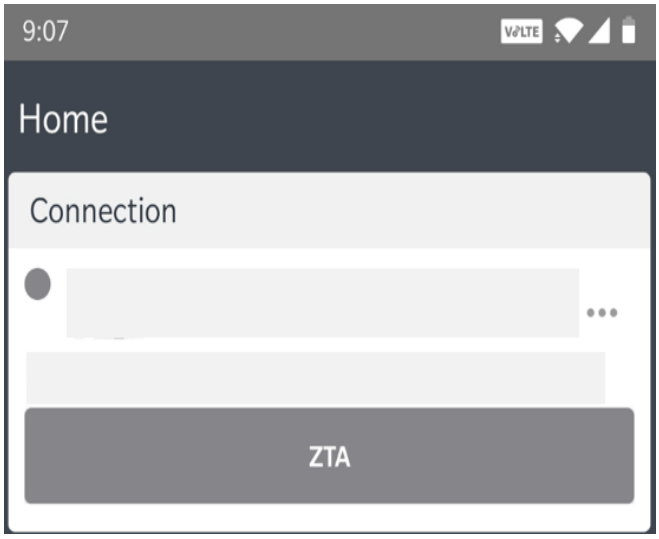

Restarting a *nZTA* connection through the *Ivanti Secure Access Client* application (Android device example)

This method also works for macOS and Windows devices.

If the existing session cookie is still valid, the *control channel* is re-established. If the session is now invalid, *Ivanti Secure Access Client* prompts the user for their *nZTA* credentials as normal. On successful re-establishment of the *nZTA* session, the user is presented with the *nZTA* End User Portal in the default browser.

When restarting *Ivanti Secure Access Client*, *nZTA* connections default to being on-demand services. That is, a previously disabled *nZTA* connection is re-enabled when *Ivanti Secure Access Client* starts.

Dynamic Policy Update and CARTA

To complement the zero-trust approach, *nZTA* supports dynamic policy updates and CARTA (Continuous Adaptive Risk and Trust Assessment) for your end user devices. This framework establishes an approach of continuous assessment and updating of secure access policies on the client, without the requirement to disconnect and reconnect to establish an updated authorization posture.

As your policies, applications, and authentication configuration are updated by the administrator on the *Controller*, changes are synchronized out to client devices dynamically and take effect immediately. *Ivanti Secure Access Client* ensures that any application updates are applied and any new authentication requirements are met before continuing the session, providing the end user with a seamlessly-updated experience. This method ensures that *Ivanti Secure Access Client* is always updated at the point of change, and not just when establishing a connection to a *nZTA Gateway* to access an affected resource.

The CARTA implementation in *Ivanti Secure Access Client* means that the security posture of the end user is continuously assessed in conjunction with policies configured in the *Controller*, with *allow* or *deny* decisions enforced through dynamic assessment and updating of the current policy set. Where application access is denied or restricted, *Ivanti Secure Access Client* informs the user of any access restrictions or policy contravention at the point of use. For example, the *nZTA* Applications Home page updates to provide visual cues with applicable error messages whenever a specific application becomes unavailable:

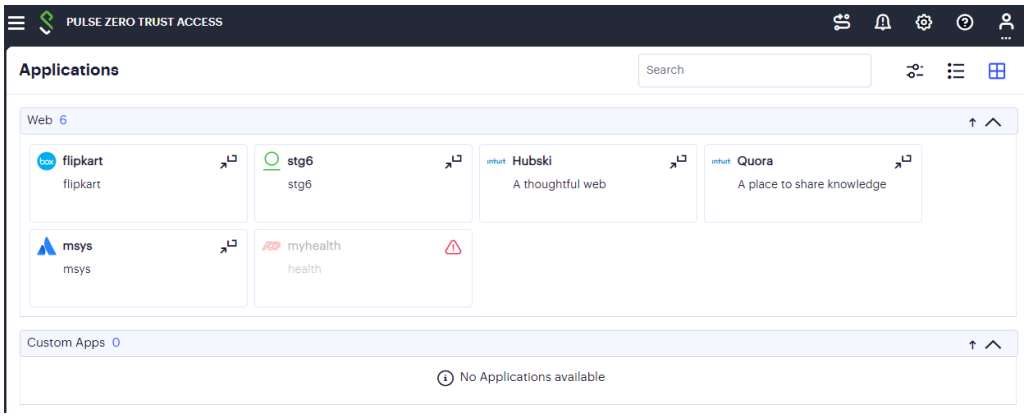

Access restricted to the "myhealth" application

By hovering your pointer over the *warning* symbol in the inactive application, *nZTA* provides an explanatory message.

Furthermore, user attempts to access a restricted web resource in a browser trigger a CARTA response with *Ivanti Secure Access Client* presenting a pop-up *resource blocked* message:

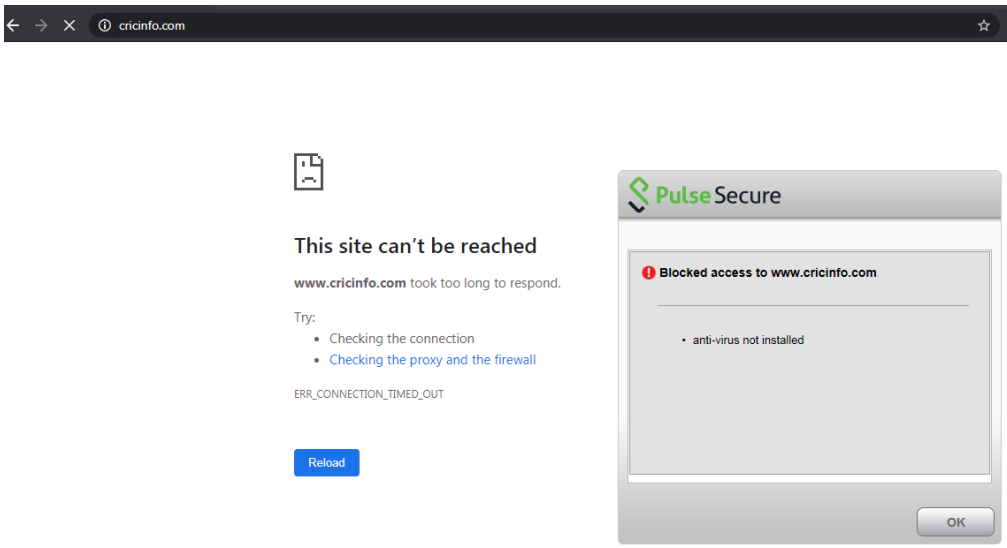

Ivanti Secure Access Client prevents browser access to a particular web resource where a device policy has been breached.

Pop-up *resource blocked* messages are not currently supported on Linux clients.

Ivanti Secure Access Client implements a no-repeat interval for *resource blocked* messages of 2 minutes, to avoid a user repeatedly seeing the same pop-up message for every browser request for the same restricted resource. While the resource remains blocked to further access attempts, no further messages are displayed by *Ivanti Secure Access Client* until after 2 minutes has elapsed. You can force *Ivanti Secure Access Client* to continue hiding *blocked resource* messages indefinitely by right-clicking the connection in the *Ivanti Secure Access Client* dialog and selecting **Disable Block Messages**. To re-enable showing *blocked resource* messages, select **Enable Block Messages**.

Using Application Discovery with *Ivanti Secure Access Client*

nZTA directs requests for a specific application towards the *nZTA Gateway* that is defined in the secure access policy for the application. All other applications use any available communications channels on the device, which may not be secure.

nZTA includes a built-in *application discovery* secure access policy on the *Controller*. This feature enables requests from any application that is not covered by a defined secure access policy to be handled by a default *nZTA Gateway*. This also enables packet analysis to be conducted from the *nZTA Gateway* to assess the validity of the requests.

To learn more about configuring default *nZTA Gateway*s and application discovery, see Configuring a Default Gateway for Application Discovery.

Client version 21.1 is required to work with application discovery and a default Gateway.

Currently, a default Gateway can only be utilized by requests from applications on macOS and Windows desktop devices.

To use application discovery on an enrolled macOS/Windows desktop device:

1. Start the *Ivanti Secure Access Client* app on the device.
2. Confirm the update of policy definitions if required.
3. Ensure that there are no other VPN connections running on the client.

   Simultaneous dual VPN connections are supported by *nZTA*, but this scenario does not support the use of a default Gateway, see On-Demand and Simultaneous Connection Handling.

4. Start the *nZTA* connection and log in with your credentials.

If another VPN connection is already in operation when the user connects to *nZTA* from the client, the default gateway configuration is ignored, and application discovery is not supported on the device. The existing VPN will handle all unassigned application requests from the client device. When a second VPN handling unassigned requests is subsequently closed on a desktop client device, use of the default gateway is enabled automatically to support application discovery on the client device. The default gateway will then handle all unassigned requests from the client device.

The *nZTA* connection starts, and available applications appear. Each listed application will use the gateway assigned by a secure access policy pushed from the *Controller*.

The default Gateway then handles all requests from applications on the macOS/Windows desktop device that are not referenced by any other secure access policy.

Using an Existing Enterprise PKI

This feature is also known as Bring Your Own Certificate (BYOC)

The *Controller* typically generates device certificates based on a trusted Certificate Authority (CA) for your end-user devices at the point of enrollment. The certificate chain and CA is installed on the device at enrollment and then the certificate is submitted for validation upon each connection to your *nZTA* service. *Ivanti* assumes that by default a tenant deployment is going to use this validation infrastructure.

However, an enterprise may wish to use its own Public Key Infrastructure (PKI) with a *nZTA* deployment where, for example, a security policy may require the use of internally-generated certificates. In this scenario, *Ivanti* enables customers to issue and manage their own end-user device certificates (for example, through a Mobile Device Management type application), and to instruct the *Controller* to use an enterprise-provided CA when validating user sessions.

Use of BYOC is subject to the following limitations:

- Support is limited to Windows and macOS devices only, using *Ivanti Secure Access Client* version 9.1R13-12191 and later.
- Gateways are restricted to using only *nZTA*-generated certificates.
- Each end-user device must be provided with its own unique certificate. *nZTA* does not support certificate reuse across multiple devices.
- Browser-based enrollment is not supported in this scenario. Users must enroll devices by creating a new *nZTA* connection from within *Ivanti Secure Access Client*.
- If a certificate expires or is revoked, an affected client device must be un-enrolled and then re-enrolled before a replacement certificate can be used.
- If you plan to add a custom domain to your subscription (see Specifying a Custom Domain), make sure you add all required certificates before you configure the custom domain.

To use your own PKI with *nZTA*, you must:

1. Request a BYOC subscription from *Ivanti*.
2. Share the required CA certificates with the *Ivanti* DevOps organization, to allow them to configure your tenants accordingly.
3. Publish your device and CA certificates to your end-user devices.

For more information, contact your support representative.

Upgrading *Ivanti Secure Access Client*

Introduction

After a desktop/mobile client device is enrolled on *Ivanti Neurons for Zero Trust Access* (*nZTA*), the following upgrade operations are supported:

- A client update package can be downloaded to each enrolled device to add new features and to resolve issues in previous versions. See Working with Client Packages.
- An Endpoint Security Assessment Plug-In (ESAP) update can be downloaded to each enrolled device. An ESAP package is included in every system software package. However, *Ivanti* releases ESAP update packages more frequently than system software versions. You may choose to upgrade the ESAP package more regularly than client software. See Enabling Minimum Supported Client Version.

All packages are automatically uploaded to the *Controller*. No manual intervention is required to achieve this. You can choose to either the latest updates to be automatically made available to your end-user client devices, or to manually control the roll out of a selected package version.

> When you select a newly-available update package for your end-users, or where you have elected to enable automatic updates and a new version becomes available, allow up to 30 minutes before the new package download is triggered on your end user client devices.

Working with Client Packages

After a client device is enrolled, you can upgrade its client software with a newer *Ivanti Secure Access Client* package from the *Controller*.

- To upgrade *Ivanti Secure Access Client* automatically, see Upgrading Ivanti Secure Access Client Automatically.

  Automatic client upgrades are not currently supported on *Ivanti Secure Access Client* Linux variants.

- To enable manual upgrade of *Ivanti Secure Access Client*, see Enabling Manual Upgrade of Ivanti Secure Access Client.

Client upgrade packages are automatically uploaded to the *Controller*. No manual intervention is required to achieve this.

Upgrading *Ivanti Secure Access Client* Automatically

Automatic client upgrades are not currently supported on *Ivanti Secure Access Client* Linux variants.

To upgrade all *Ivanti Secure Access Client* to the latest client package automatically:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Administration** icon, then select **Upgrade > Installation Packages**.

   The *Installation Packages* page appears.
3. Click the **Ivanti Secure Access Client** tab.

   A list of available *Ivanti Secure Access Client* packages appears.
4. Enable the **Auto update to latest version** control.

   The most recent *Ivanti Secure Access Client* package is selected automatically. For example:

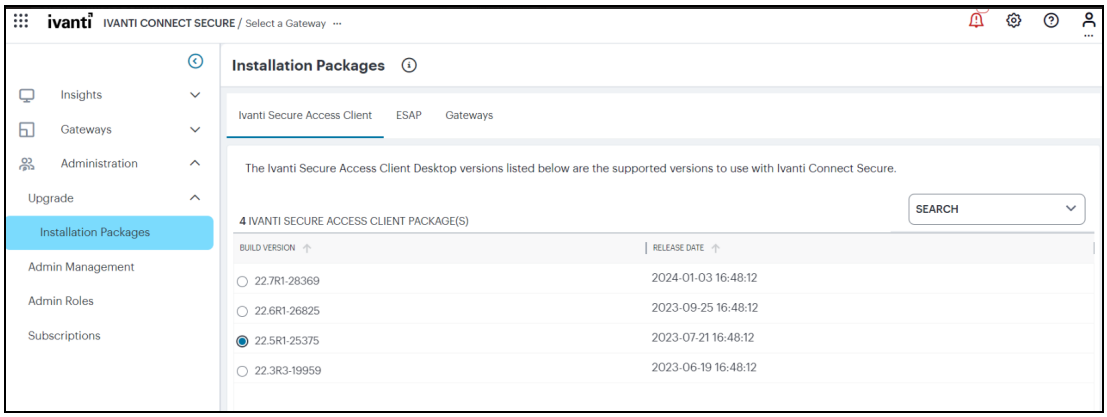

Automatic *Ivanti Secure Access Client* Update

As each user next logs into *Ivanti Secure Access Client* on their device, if their software is at an earlier version than the latest version in the *Controller*, *Ivanti Secure Access Client* provides an automatic update to the selected version.

Enabling Manual Upgrade of *Ivanti Secure Access Client*

To enable users to manually update to a specified version of *Ivanti Secure Access Client*:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Administration** icon, then select **Upgrade > Installation Packages**.

   The *Installation Packages* page appears.
3. Click the **Ivanti Secure Access Client** tab.

   A list of available *Ivanti Secure Access Client* packages appears.
4. Disable the **Auto update to latest version** control.
5. Select any of the listed *Ivanti Secure Access Client* packages. This is the version of the *Ivanti Secure Access Client* software that you want users to have on their device. For example:

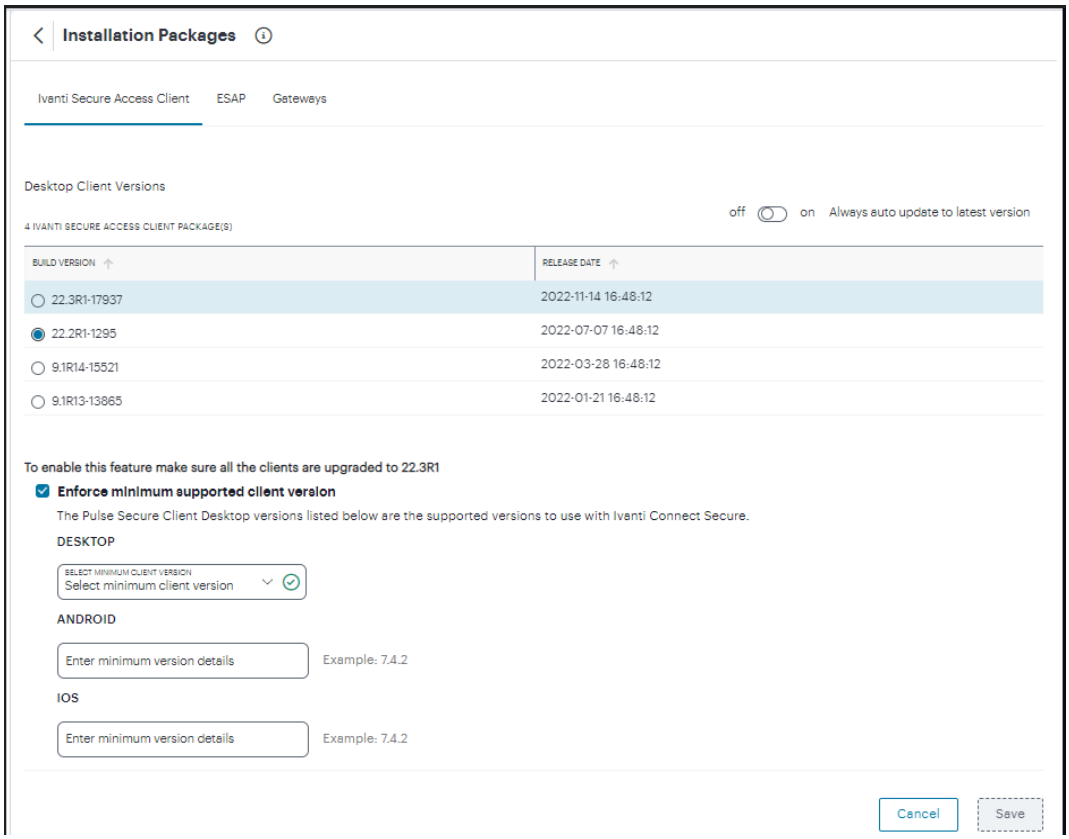

Manual *Ivanti Secure Access Client* Update

As each user next logs into *Ivanti Secure Access Client* on their device, if their software is at a different version, *Ivanti Secure Access Client* provides a prompt to the user to change to the version you selected in the *Controller*.

Ivanti Secure Access Client for Desktops does not support a downgrade path. If the version you select in the *Controller* is earlier than the client version currently installed on a user's desktop device, the user must first uninstall the current *Ivanti Secure Access Client* software. To trigger installation of the required version, the user must then re-perform the *nZTA* enrollment procedure through a web browser.

Enabling Minimum Supported Client Version

This feature provides an option for the admin to force upgrade the client version when Minimum Client Version is configured.

To enable this feature:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Administration** icon, then select **Upgrade > Installations**.

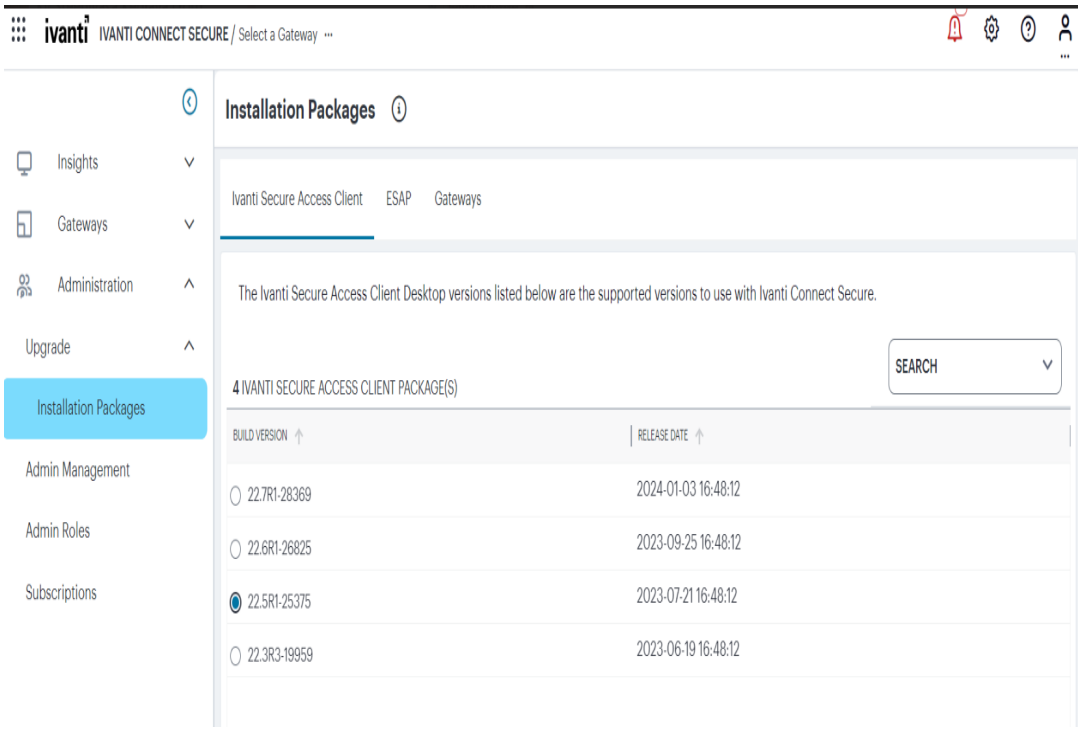

Enforce Minimum Client Version

3. Select the **Enforce minimum client version** check box.

   The following options appear.

   - Desktop
   - Android
   - iOS

   If **Force update** is selected, Client will auto-upgrade once you enroll with older client.

4. Specify the minimum client version to be enforced on Ivanti Secure Access Client. For example, 22.3R1.

   If the minimum client version is not specified for any specific type of client, any version of the client is allowed to connect to Ivanti Connect Secure without any minimum client version enforcement for that specific client.

5. Click **Save**.

Working with ESAP Packages

ESAP (Endpoint Security Assessment Plug-in) is a *Ivanti Secure Access Client* plug-in through which you can upload the latest security assessment definitions independently of a full client software upgrade. After a user device is enrolled and *Ivanti Secure Access Client* has been installed, you can upgrade ESAP with newer versions from the *Controller*.

- To update ESAP definitions on *Ivanti Secure Access Client* automatically, see Updating ESAP Definitions on Clients Automatically.
- To enable the manual update of ESAP definitions on *Ivanti Secure Access Client*, see Enabling Manual Update of ESAP Definitions on Clients.

> ESAP upgrade packages are automatically uploaded to the *Controller*. When each user next logs into *Ivanti Secure Access Client* on their device, if their ESAP definition is at an earlier version, the ESAP definition updates automatically from the most recent ESAP package.

Updating ESAP Definitions on Clients Automatically

To view all available ESAP packages:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Administration** icon, then select **Upgrade > Installation Packages**.

   The *Installation Packages* page appears.
3. Click the **ESAP** tab.

A list of available ESAP packages appears.

> This feature is permanently enabled. When each user logs into *Ivanti Secure Access Client* on their device, if their ESAP definition is at an earlier version, the ESAP definition updates automatically from the most recent ESAP package.

Enabling Manual Update of ESAP Definitions on Clients

> This feature is not currently supported.

Using the Insights Menu to Monitor User Activity and Service Usage

Introduction

Ivanti Neurons for Zero Trust Access (*nZTA*) provides visibility of user activity and service usage across your enterprise through network activity analytics, gateway performance graphs, application usage metrics, and stored activity logs.

After you log in to the Tenant Admin Portal following successful completion of the *Onboarding Wizard*, *nZTA* displays the **Network Overview** page. This page presents a top-down overview of your application infrastructure, providing an opportunity to monitor user and *nZTA Gateway* activity, and to identify problems and compliance issues as they occur. For more information, see Reviewing Your Network Activity.

Through the *nZTA* menu, use the **Insights** menu icon to:

- View graphs, metrics, and logs concerning user activity, see Reviewing User Activity.
- See details and usage of applications configured in your *nZTA* service, see Reviewing Application Usage.
- Monitor *nZTA Gateway* activity, see Monitoring nZTA Gateway Activity.
- Review policy failures, see Reviewing Policy Failures.
- View activity logs, see Checking the Logs.
- Configure actionable insights, see Actions.
- Obtain reports of activity and usage across your services, see Reports.
- View events and notifications, see Viewing Alerts and Notifications.

nZTA provides both a light theme and a dark theme for the UI display. To learn more, see Changing the UI Theme.

No analytics data will be displayed on any dashboards when nZTA Gateway is bypassed.

Using Filters and Selectors to Monitor Specific Services

Each page in the *Insights* menu allows data filtering through the filter bar (see Using the Filter Bar), enabling you to observe and monitor only the analytics and log data you want. Filters fall broadly into two categories, and are applied as applicable to the page you are viewing:

- *Summary page filters*: high-level filters and selectors such as time period and Gateway, user, or application, that apply across both summary and detail *insights* pages. Filters applied here can affect the data on all *Insights* pages that you visit. For example, the same selected time period remains in place across every page.
- *Detail page filters*: filters that are applied at a chart detail page view that are applied to the log data constituting the chart being interrogated. For more details, see Viewing Detailed Logs for a Chart.

  The **Logs** page uses a separate time period selection filter from other *Insights* pages. A time period selected here is not applied elsewhere, and vice versa, yet is retained across login sessions in the same way.

Filters and selection criteria are persistent across all relevant *Insights* pages, and are retained across login sessions. When you log back in, the same selection criteria remain applied. Settings are stored per admin user, such that each admin maintains their own view of the analytics data.

You can remove applied filters and return to the default setting through the *reset* option in most filter controls. For example:

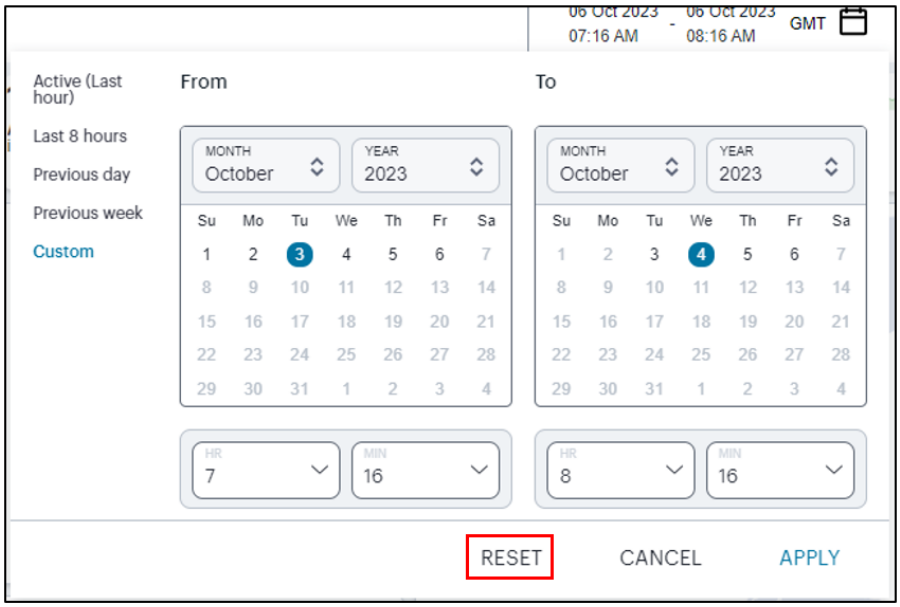

Resetting your selected time period filter (indicated) back to the default "Last Hour" active data view

Or in the case of each chart detail or log page, the log filter bar typically includes a reset icon. For example:

Using the Reset icon (indicated) to reset log filters to their default state

Information on the filters and data selection options available on each page is provided throughout this guide. Refer to the page-specific help and documentation for more details.

> To perform a global reset of all applied filters and selection criteria on all pages, in a single action, use the **Settings** menu *Reset Filters* option. To learn more, see Resetting All Filters and Selections.

Reviewing Your Network Activity

The *Network Overview* page shows real-time analytics data for your application infrastructure, providing a one-page dashboard of activity across your organization.

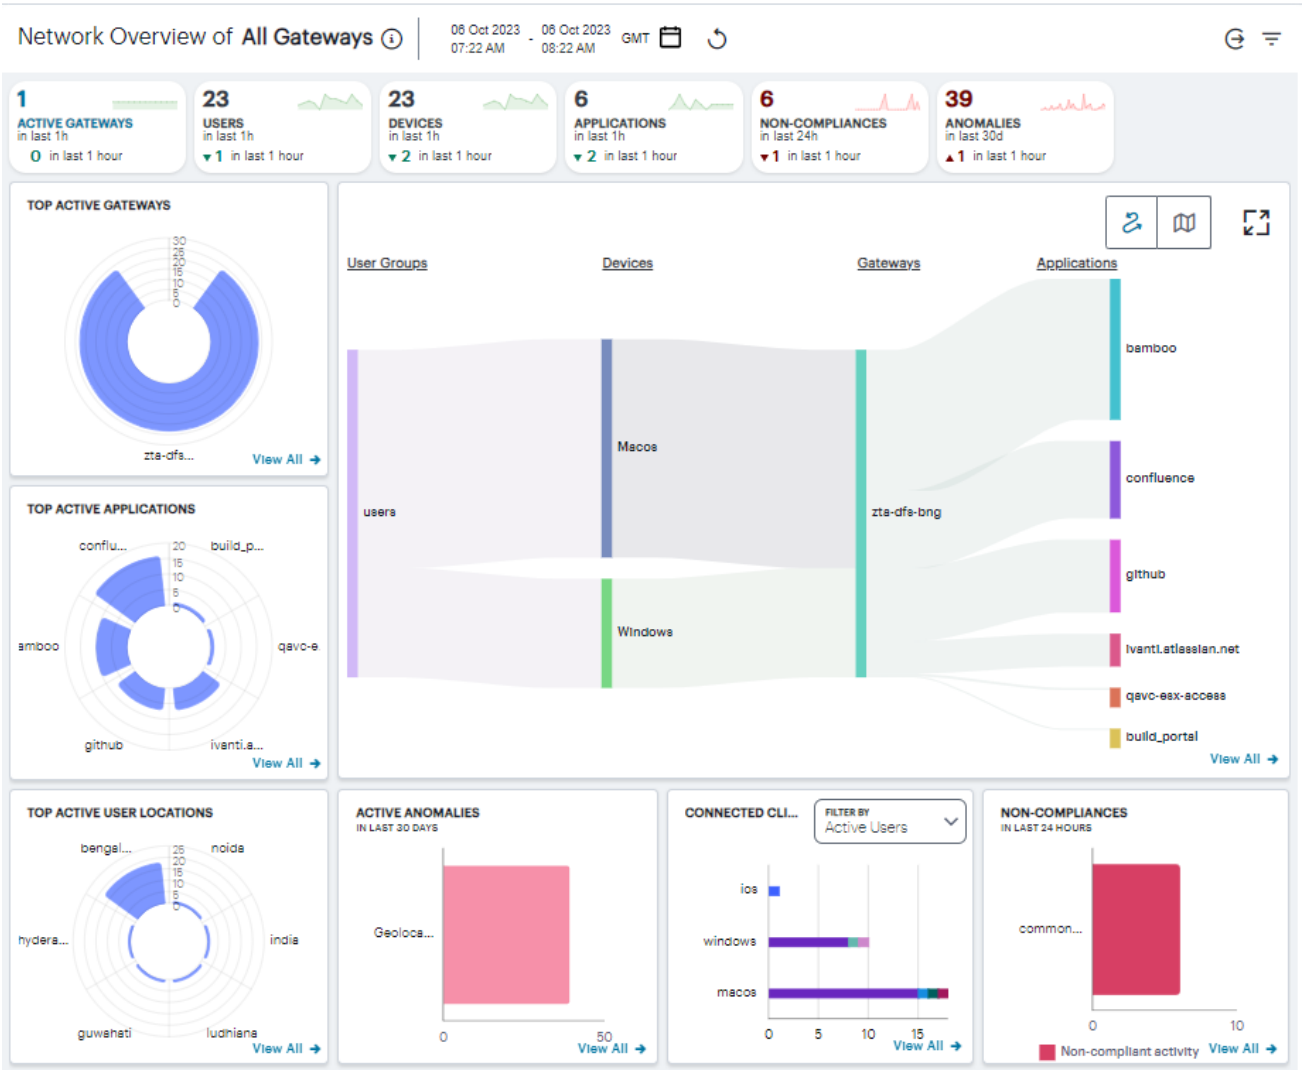

An overview of Network activity across your enterprise

To access the **Network Overview** page:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.

   The *Network Overview* page appears by default.

2. To return to the *Network Overview* page at any time, select **Insights > Overview** from the left-hand menu. Alternatively, select the *Ivanti* banner at the top.

Understanding the Display

The primary components of the **Network Overview** page are the following:

- **Filter bar**, allowing the selection of active or historic data. For details, see Using the Filter Bar.
- **Summary ribbon**, showing totals for active **Gateways**, **Users**, **Devices**, **Applications**, **Compliance Failures**, and **Anomalies**. For more details, see Using the Summary Ribbon.
- **Switchable World Map and Sankey Chart views**, showing active Gateway or user locations. Sankey chart view is the default view.
    - In the world map view, each indicated location provides a summary of the activity observed there. For more details, see Using the World Map.
    - In the Sankey chart view, you can view the relationships between user groups, device types, *nZTA Gateway*s, and applications. For more details, see Using the Sankey Chart View.
- **Radar charts**, providing top usage data for **Gateways**, **Applications**, and **User Locations**. For more details, see Using the Top Active Breakdown Charts.
- **Bar chart** breakdowns showing **Active Anomalies**, **Connected Clients Version**, and **Non-compliance** activity. For more details, see "Using the Active Anomaly, Connected Clients Version, and Non-Compliance Charts" on page 613.

The data in this page refreshes automatically every 5 minutes.

With each chart, click the **View all** link to view a page of detailed log records for that category. To learn more about using the chart detail page, see Viewing Detailed Logs for a Chart.

The following principles apply to all elements of the page:

- A user can have one or more devices.
- Each device can have only one active secure access session.
- One session can connect to multiple applications.
- One session can be associated with multiple *nZTA Gateway*s.
- One *nZTA Gateway* can have multiple applications registered with it.
- One application instance can be registered with only one *nZTA Gateway*.

Using the Filter Bar

nZTA uses the top part of the display on all **Insights** data analysis pages to show the current page title, the selected time period and timezone, and options to:

- Select the date and time period for which data is displayed
- Manually refresh the data
- View analytics data for a selected user or application
- Filter analytics data by a selected *nZTA Gateway*

To learn more about how filters are applied in the Tenant Admin Portal, see Using Filters and Selectors to Monitor Specific Services.

> To configure the default timezone for the data displayed in this admin login account, see Setting the Timezone.

By default, analytics data on all pages is shown for the last hour. To select a previous or specified time period, select the date-time display (indicated):

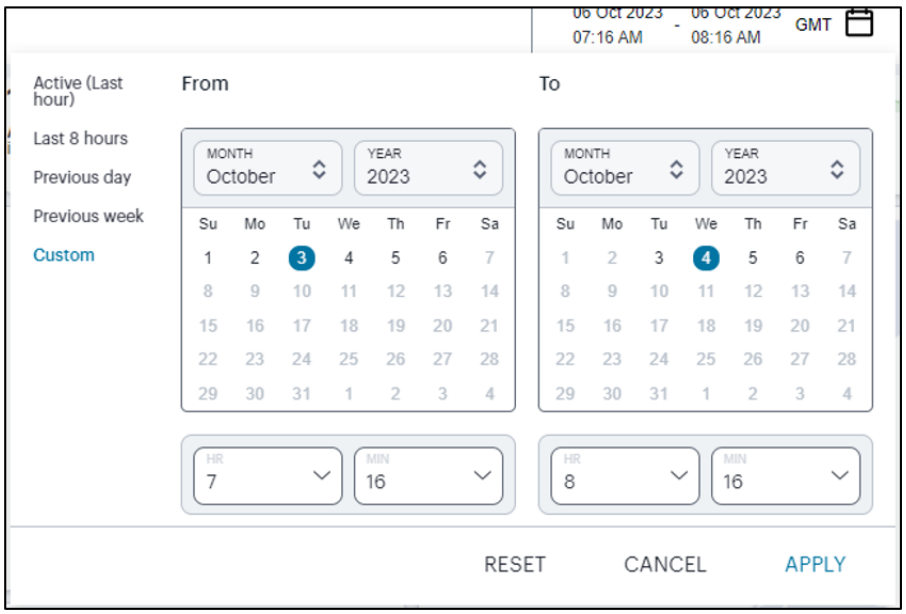

Selecting a date and time range

In the date-time selection dialog, choose from the following predefined time period options:

- **Last hour**: Data observed for the previous 60 minutes.
- **Last <X> hours**: Data observed so far in the current day, up to the last hour (in GMT).
- **Previous day**: Data observed for the previous full day.
- **Previous Week**: Data observed for the previous calendar week (for the previous full Sunday-to-Saturday week).
- **Custom**: Data observed for a chosen time period. If you select this option, *nZTA* enables you to select a custom time period using the **From** and **To** date/time calendar controls.

> The date/time calendar controls are enabled for only the **Custom** option. However, the calendar continues to identify the applicable start and end date-time for all predefined time periods.

To reset the selected time period back to the default (*Last Hour*) view, select **Reset**. To return to the current page without making any changes, select **Cancel**.

To apply your changes, select **Apply**. The selected time period is displayed in the filter bar and data across all **Insights** pages is updated accordingly.

The data in the display refreshes automatically at 5 minute intervals. To manually refresh the data, click the circular arrow:

Refreshing the data

nZTA provides the ability to show focused metrics for individual users or applications. To select a specific user or application, use the following icon:

Selecting a specific user or application

Then, from the drop-down menu, select one of the available options:

- Select **Set User** to view data for a selected user. In the search box provided, start typing a user ID. *nZTA* auto-completes any matching user IDs. Next, select **View User**.

  The *User Activity* page appears. To learn more, see Showing Activity for a Specific User.

- Select **Set Application** to view usage metrics for a selected application. In the search box provided, start typing an application name. *nZTA* auto-completes any matching names. Next, select **View Application**.

The *Application* page appears. To learn more, see Showing Usage Data for a Specific Application.

> You can also access data for individual users or applications by selecting the name of a user or application from the corresponding info-panel, activated through the Summary Ribbon. For more details, see Using the Summary Ribbon.

nZTA also provides the ability to set a Gateway filter on all *Network Overview*, *User*, *Application*, and *Policy Failure* analytics pages in the *Insights* section. Applying a Gateway filter means that all dashboards are updated to show only activity relating to the chosen *nZTA Gateway*. In other words, *nZTA* shows only analytics for applications that were accessed from that specific *nZTA Gateway*, along with *nZTA Gateway* activity for users and devices being active in the selected time period. This filter is persistent across all pages, and remains in place for the duration of the current login session.

By default, the filter is inactive and shows data for *All Gateways*, as indicated in the title of all pages.

To set a Gateway filter, select the Gateway filter icon:

Filtering analytics data by *nZTA Gateway*

Then, from the Gateway selection panel, choose your *nZTA Gateway* from the drop-down list:

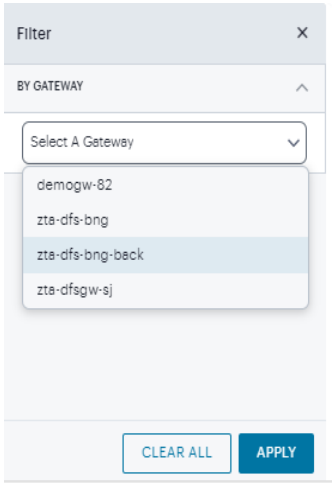

Using the Gateway filter panel

To set the Gateway filter, select **Apply**.

To remove a filter and return to viewing analytics for *all gateways*, select **Clear All**.

> On detailed log pages for individual charts (see Viewing Detailed Logs for a Chart.) you cannot set a Gateway filter directly. Instead, set the Gateway filter on the parent page before you click through to the individual chart logs.

Using the Summary Ribbon

The Summary Ribbon at the top of the page shows data totals for the selected time filter:

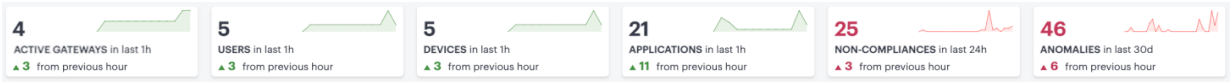

Viewing the summary ribbon

The ribbon indicates the totals accrued for each category during the displayed time period, as indicated adjacent to the category name.

The following categories are provided in the ribbon:

- The number of **Active Gateways**.
- The number of active **Users**.
- The number of active **Devices**.
- The number of in-use **Applications**.
- The number of **Non-compliances**. In other words, non-compliant attempts to access your applications. For the default time period filter, non-compliance totals shown here are for 24 hours. For other selected time periods, the number reflect the total for that period.
- The number of **Anomalies** detected by *nZTA*. That is, the total number of geographic and business hours anomalies. For the default time period filter, anomaly totals shown here are for the previous 30 days, and include only unacknowledged anomalies. For other selected time periods, this total includes both acknowledged and unacknowledged anomalies.

Compliance and anomaly counters use the following color scheme to reflect status:

- **Black**: No geographic anomalies or compliance failures are reported
- **Red**:
    - **Non-compliance**: if the count is non-zero
    - **Anomalies**: if the count is non-zero

If you are currently viewing data for the *last hour*, each category in the ribbon includes a trend graph (highlighted, top) showing the changes in data during the hour. Also included is a change value (highlighted, bottom) based on the previous hour:

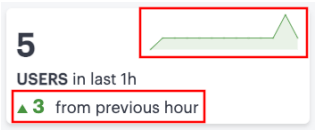

Data trends for this hour versus the previous hour

> In the default *last hour* view, while data for Active Gateways, Users, Devices, and Applications is shown as such, non-compliances are shown for the previous 24 hours and anomalies are shown for the previous 30 days. This is as indicated against the Category name.

Additional trend indicators are present for the *last hour* time period only. All other time periods show only the main data totals for each category.

If you click on any of the categories in the ribbon, *nZTA* displays a sliding info-panel dialog showing more details for that category. For example, if you click on the **Active Gateways** category, a panel appears showing the list of active *nZTA Gateway*s. In this case, a summary box is displayed for each *ZTA Gateway* showing statistics relevant to that instance, such as instance health (disk, CPU, and memory utilization), the number of active users, applications, active devices, non-compliance events.

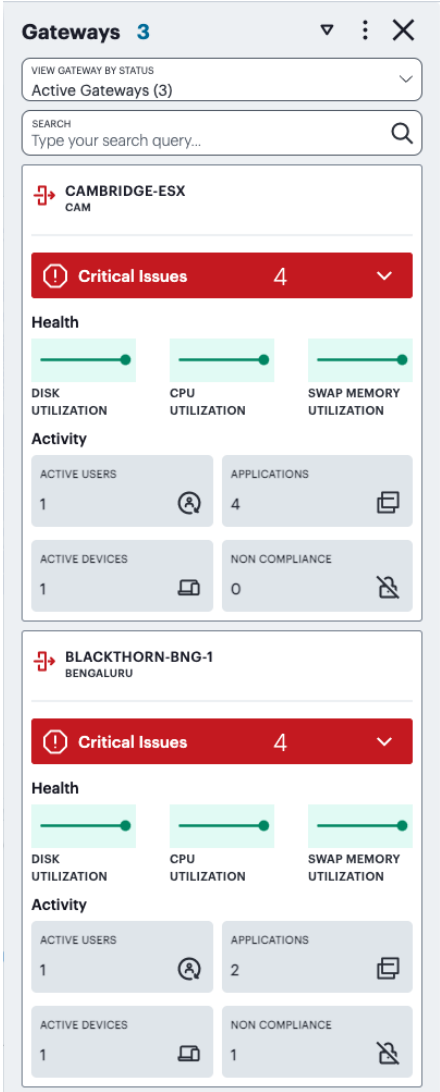

Viewing the Gateways info-panel

The following color scheme is used in the icon adjacent to the item listed in the panel:

- **Black** / **Green**: No issues are reported for the item shown in the info-panel
- **Red**:
  - **Users info-panel**: The user has anomalies reported against them in the selected duration
  - **Gateways info-panel**: The Gateway is reporting critical issues

When displaying active Gateway data, all non-compliance and unacknowledged anomaly totals are displayed for the previous 24 hours.

The *Gateways* info-panel displays the following details for each Gateway in your deployment:

- **Location name and number of Gateways**: The descriptor for this location and the number of Gateway instances deployed there.
- **Warning/Critical Issues**: A list of warnings or critical issue messages reported by the Gateways at this location.
- **Gateway Health**: Health indicators for the Gateways at this location.
- **Active Users**: The number of unique users accessing applications through Gateways at the location (as also indicated in the location counter)
- **Active Applications**: The number of applications accessed through Gateways at the location
- **Active Devices**: The number of unique devices used to access applications through Gateways at the location
- **Non-Compliant**: The number of non-compliant access attempts to applications configured for Gateways at the location (note that attempts by the same device to access two applications for which it does not meet compliance requirements increment this total by two)

This version of the info-panel shows details for *all Gateway locations*. To view an info-panel for a single *nZTA Gateway* location, click the Gateway location counter in the world map. For more details, see Using the World Map.

Use the **View Gateway by Status** drop-down list to change the type of Gateways displayed in the panel. Choose from:

- **All Gateways**: All *nZTA Gateway*s regardless of status.
- **Active Gateways**: All active *nZTA Gateway*s. That is, only those *nZTA Gateway*s that are responsive, irrespective of health status, and have observed application accesses during the selected time period. This is the default view.
- **Offline Gateways**: All offline *nZTA Gateway*s. That is, only those *nZTA Gateway*s that are unresponsive.
- **Online Gateways**: All online *nZTA Gateway*s. That is, only those *nZTA Gateway*s that are responsive but have not observed any application accesses.
- **Unregistered Gateways**: All currently unregistered *nZTA Gateway*s. That is, only those *nZTA Gateway*s that are deployed but not yet registered with the *Controller*.

The number of instances of each type is given in brackets.

For example, by selecting *Offline Gateways*, the panel updates as follows:

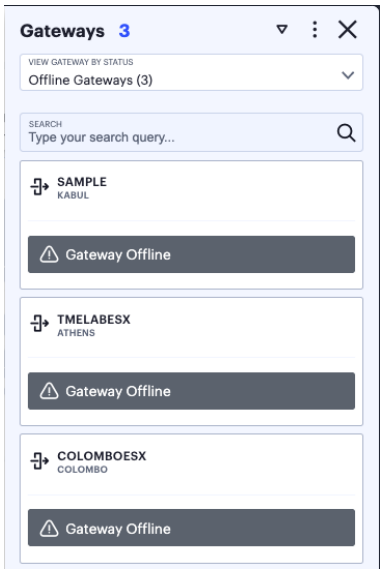

Viewing all offline *nZTA Gateway*s in the Gateways info-panel

Use the **Search** bar at the top to filter the results list. For example, to show only those *nZTA Gateway*s that match a search string. To clear your search, click **CLEAR SEARCH RESULTS**.

Hover your pointer over the instance health indicators to display a tooltip showing more specific details and values.

Click on any Critical or Warning notification banner to display a drop-down summary of the issues:

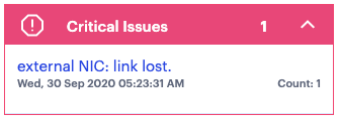

Viewing critical issues

You can click on each entry to obtain more details and logs concerning the issue.

For the **Active Users** info-panel, *nZTA* displays an *average UEBA Threat score*. To learn more about UEBA Threat scores, see Showing Activity for a Specific User.

For Non-Compliance and Anomalies info-panels, summaries are displayed on a per-user basis, with the reason for the event shown.

To change the sort order of the items displayed in the info-panel, use the *Sorting* controls at the top:

Changing the info-panel sort order

Use the *dots* icon to select the sort criteria, then use the *arrow* icon to toggle between ascending and descending order. The sort criteria varies depending on the category chosen, and is based on the statistics shown for each item. For example, by selecting the **Gateways** info-panel, you can choose the display order for your *nZTA Gateway*s based on the following statistics:

- Active Users
- Apps Accessed
- Non-compliances
- Active Devices
- Number of Issues
- Gateway Name
- City Name

A tick identifies the currently chosen criteria.

For **Anomalies**, the info-panel provides additional functionality to enable you to:

- Acknowledge individual anomalies and remove them from the active total.
- Filter on acknowledged, unacknowledged (active), or all anomalies.
- Terminate the corresponding active user session, if applicable.

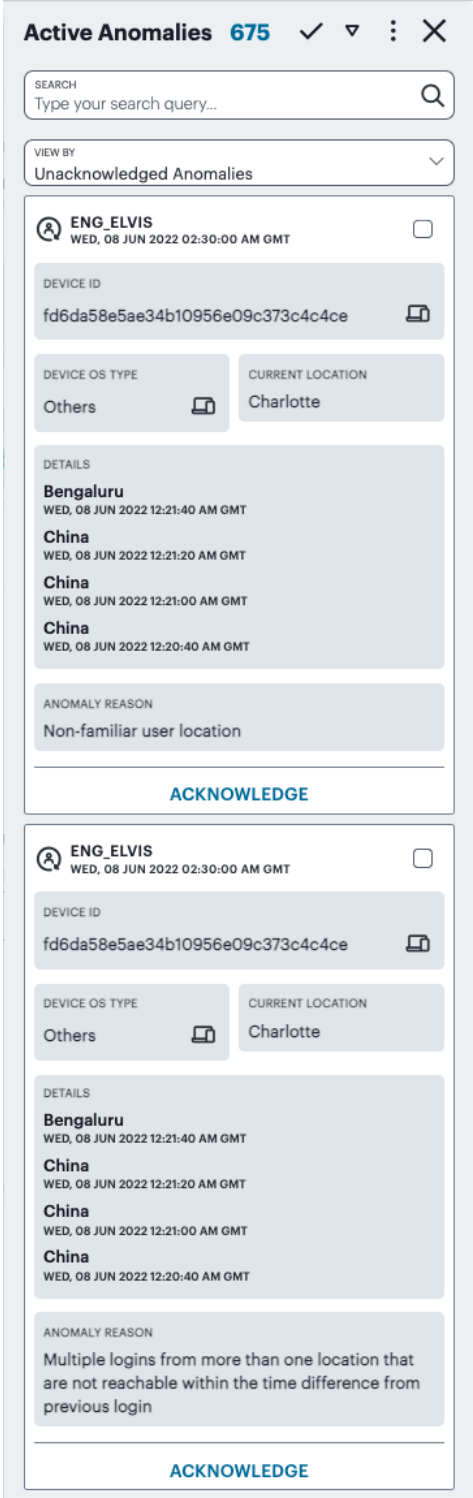

Viewing the Anomalies info panel

Each box in the info-panel lists a user and the active anomalies connected to them. For each user, click **ACKNOWLEDGE** to remove this anomaly from the list. Alternatively, use the *tick* icon and check boxes adjacent to each user name to acknowledge multiple, or all, anomalies in a single action. Note that when the default "active" time period filter is selected, the anomalies count in the summary ribbon decreases by 1 for each acknowledgment. To terminate the active session for this user with immediate effect, click **END SESSION**. Session termination is available only for admin users with full access permissions.

> For other selected time period filters, the anomalies total includes both acknowledged and unacknowledged anomalies.

To view (and optionally terminate) sessions for all active users, see Viewing and Terminating User Sessions.

For each version of the info-panel, you can click the name of an item listed in the panel to access further pages that provide usage metrics or configuration details for that item:

For **Non-compliances**, the info-panel provides additional functionality to enable you to:

- Acknowledge individual non-compliance and remove them from the active total.
- Filter on acknowledged, unacknowledged (active), or all non-compliances.

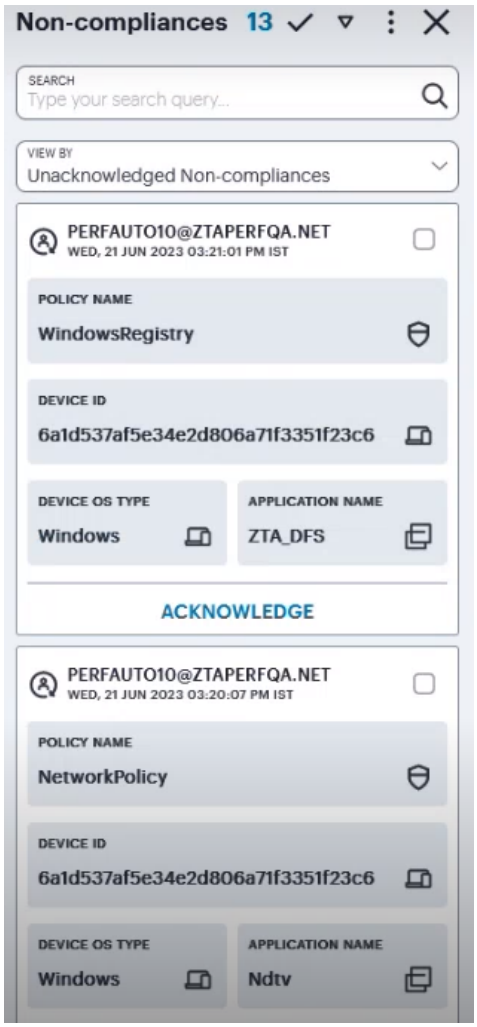

Viewing the Non-compliances info panel

Each box in the info-panel lists a user and the active non-compliances connected to them. For each user, click **ACKNOWLEDGE** to remove this non-compliance from the list. Alternatively, use the *tick* icon and check boxes adjacent to each user name to acknowledge multiple, or all, non-compliances in a single action. Note that when the default "active" time period filter is selected, the non-compliances count in the summary ribbon decreases by 1 for each acknowledgment.

> For other selected time period filters, the non-compliances total includes both acknowledged and unacknowledged non-compliances.

To view (and optionally terminate) sessions for all active users, see Viewing and Terminating User Sessions.

For each version of the info-panel, you can click the name of an item listed in the panel to access further pages that provide usage metrics or configuration details for that item:

- For the **Gateways** info-panel, click a *nZTA Gateway* name to access the corresponding *Gateways Overview* page. For more details, see Viewing and Monitoring Gateways in the Controller.
- For the **Users** info-panel, click a user name to access the corresponding *Users* analytics page. For more details, see Showing Activity for a Specific User.
- For the **Devices** info-panel, click a device type to access the *Device Configuration* page, filtered by that device type. For more details, see Viewing Currently Enrolled User Devices.
- For the **Applications** info-panel, click an application name to access the corresponding *Application* analytics page. For more details, see Showing Usage Data for a Specific Application.
- For the **Non-compliances** and **Anomalies** info-panels, click a user name to access the corresponding *Users* analytics page. For more details, see Showing Activity for a Specific User.

Using the World Map

The *world map* provides a geographically-positioned view of your Gateway or user locations, selected through the switcher at the top of the panel:

- Select **Gateways** (the default setting - indicated) to display your Gateway locations on the map as a series of geographically-placed counters.

Viewing Gateway locations on the world map

Each counter shows the status of the services held there and the number of active user connections. Gateway status is indicated by the color scheme shown in the legend:

- **Good (Green)**: All Gateways are functioning normally.
- **Warning (Amber)**: One or more of the Gateways at that location is experiencing a *warning* scenario. This status is triggered by the occurrence of any one of the following conditions:
    - Gateway device CPU usage is within the range 80% - 90%
    - Gateway device swap memory usage is within 10% - 50%
    - Gateway device disk usage is within the range 80% - 90%
- **Critical (Red)**: One or more of the Gateways at that location is experiencing an *critical* alert scenario. This status is triggered by the occurrence of any one of the following conditions:
    - Gateway device swap memory usage is greater than 50%
    - Gateway device disk usage is greater than 90%
    - At least 1 critical error has been reported
- **Offline (Gray)**: One or more of the Gateways at that location is offline and/or unresponsive, or is not yet registered with the *nZTA* Controller.

Hover your pointer over a counter to view a visual representation of the users currently connected to the Gateways at that location. The greater the number of users at an originating location, the larger the indicator on the map:

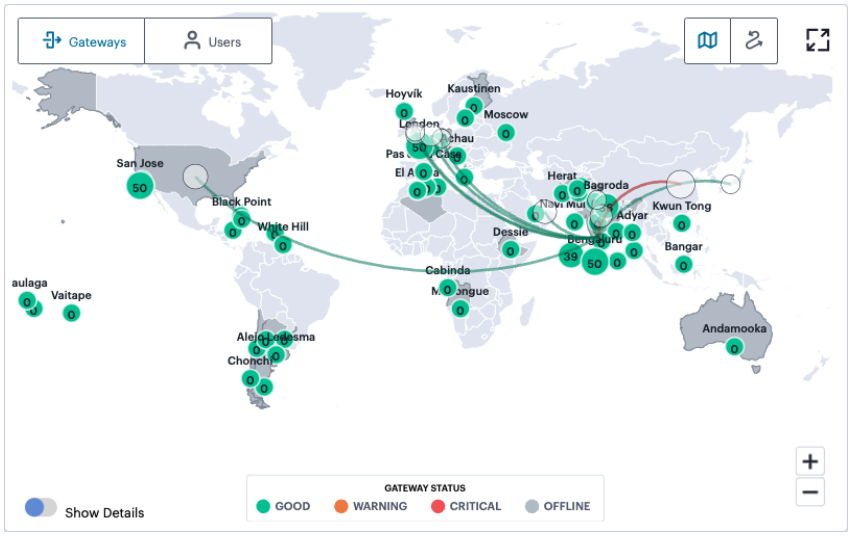

Viewing connected users for a selected Gateway

In this view, a red connecting line between a user location and a Gateway location indicates non-compliances exist for those user devices.

In addition, use the Show Details switch to toggle on or off a tooltip summary panel for the Gateway location that overlays the display:

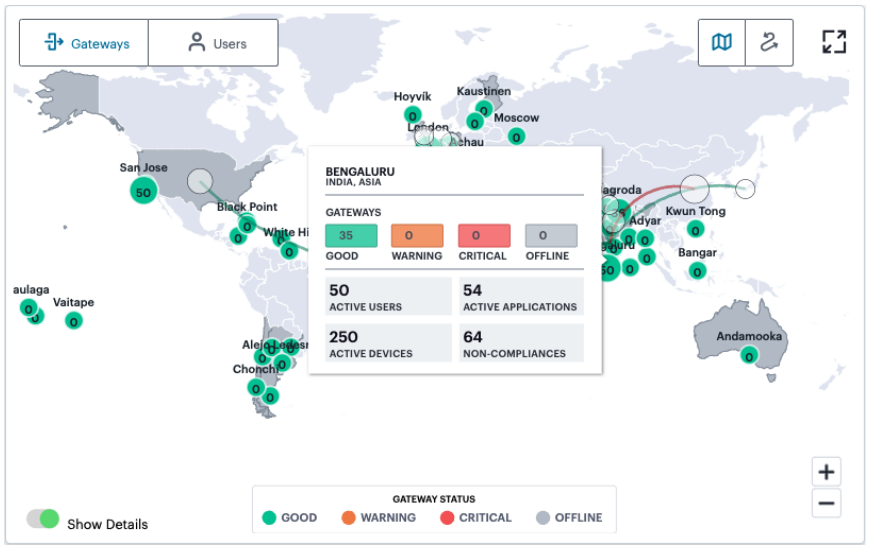

Viewing a location status tooltip panel for Gateways

This panel indicates the status of the Gateways at that location and provides metrics concerning the status of the services at that location:

- **Active Users**: The number of unique users accessing applications through Gateways at the location (as also indicated in the location counter)
- **Active Applications**: The number of applications accessed through Gateways at the location
- **Active Devices**: The number of unique devices used to access applications through Gateways at the location
- **Non-Compliances**: The number of non-compliant access attempts to applications configured for Gateways at the location (note that attempts by the same device to access two applications for which it does not meet compliance requirements increment this total by two)

Select a counter to show the Gateways info-panel for the individual location. For more information, see Using the Summary Ribbon.

This view of the info-panel displays data for a single Gateway location. To view an info-panel showing data for all Gateway locations, click the Gateways category in the Summary Ribbon. To learn more, see Using the Summary Ribbon.

- Select **Users** (indicated) to display your user locations on the map:

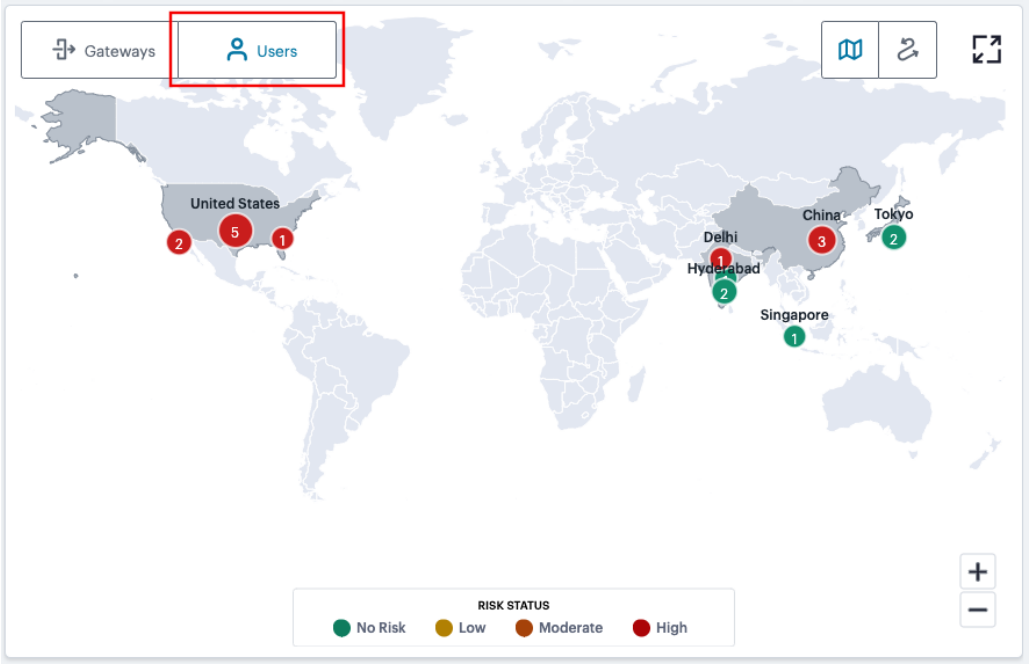

Viewing user locations on the world map

In this view, each counter shows the number of users at a geographic location that are connected to your Gateways. Gateway status is indicated by the color scheme shown in the legend:

- **No Risk (Green)**: All users at this location have a UEBA Threat score that does not exceed the threshold for zero risk.
- **Low (Amber)**: One or more users at this location have a UEBA Threat score that falls in the range defined as a low risk.
- **Moderate (Orange)**: One or more users at this location have a UEBA Threat score that falls in the range defined as a moderate risk.
- **High (Red)**: One or more users at this location have a UEBA Threat score that falls in the range defined as a high risk.

Hover your pointer over a counter to show a tooltip panel containing the UEBA Threat score summary for those users:

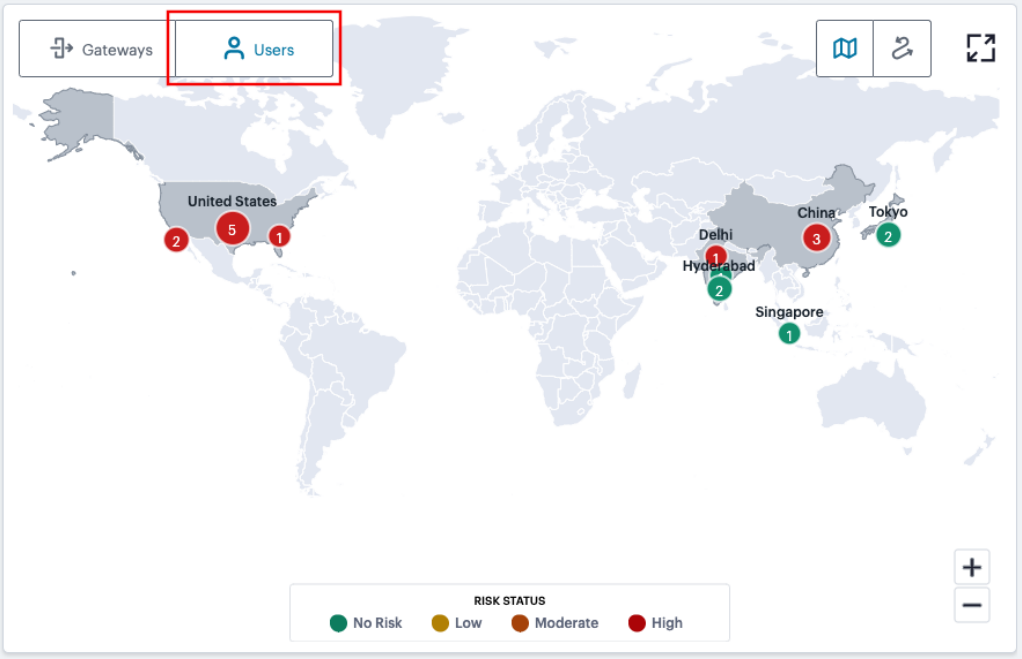

Viewing user locations on the world map

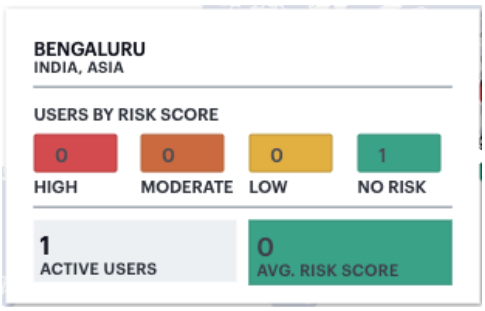

Viewing a location status tooltip panel for users

Select a counter in this view to show the *Users* info-panel. For more information, see Using the Summary Ribbon.

In both views, use the Plus (+) and Minus (-) controls to zoom in and out of the world map, allowing you to select the desired level of detail. Alternatively, use your pointer to manipulate the map display. Double-click/tap an open area of the map to zoom in, or reposition the map display through drag and drop.

To toggle between the Map view and Sankey chart view, use the icons at the top-right:

Toggle between Map view and Sankey chart view

The data shown is representative of the currently-selected time period, and by default shows *active* data (for the previous 1 hour). To learn more about setting time periods for the displayed data, see Using the Filter Bar.

To expand the current view, click the Full Screen icon:

Expand the current view

Click the Full Screen icon again to return to the standard view.

Using the Sankey Chart View

The Network Sankey chart provides an alternate visualization of your services, showing directed flow between related objects. The width of each stream in the flow is proportional to the utilization of the object the flow passes through, allowing an administrator to view significant usage and relationships across your user base and application infrastructure.

To activate the Sankey chart view, use the toggle icons at the top-right:

Toggle between Map view and Sankey chart view

By clicking the toggle display icon, the Sankey chart replaces the world map in the display. All other components remain unchanged.

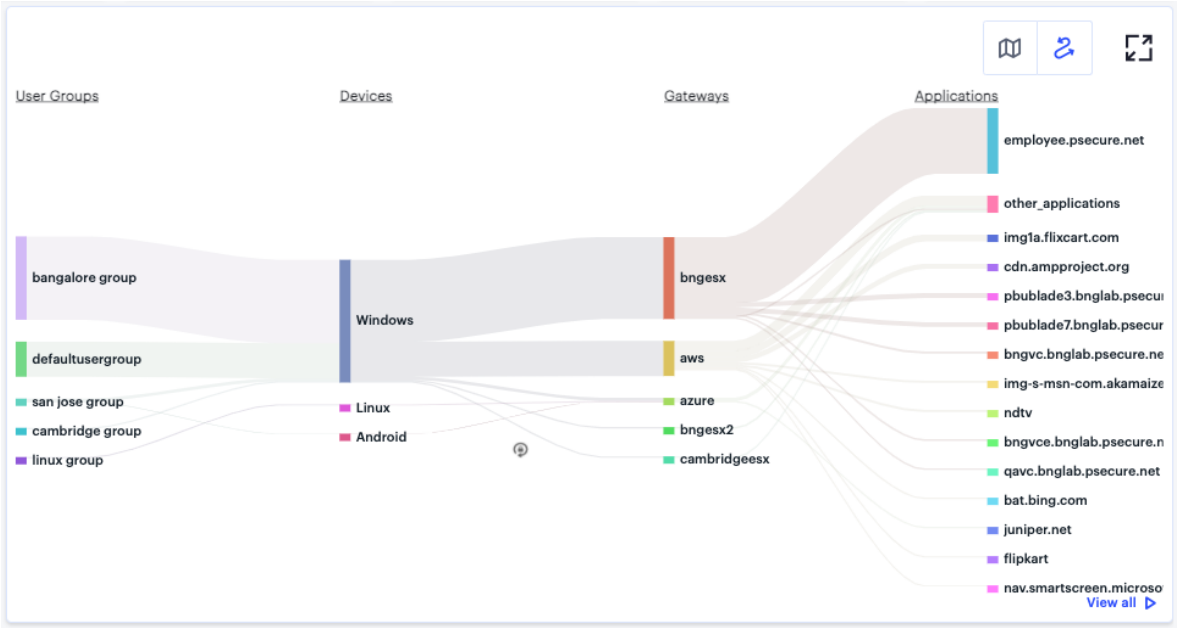

Displaying the Network Overview Sankey Chart View

The *nZTA* Sankey chart maps **User Groups** > **Device Types** > **Gateways** > **Applications**. By hovering your pointer over a flow of interest, *nZTA* displays a tooltip confirming the scale of the activity between the two objects connected by the flow.

To focus the display on a specific flow, or to identify related objects that interact with this flow, click the chart at a point of interest. *nZTA* provides highlighting to all flows that pass through the point selected.

Using the Active Anomaly, Connected Clients Version, and Non-Compliance Charts

The **Network Overview** page includes bar charts to provide a breakdown of **Active Anomalies**, **Connected Clients Version**, and **Non-compliance** events.

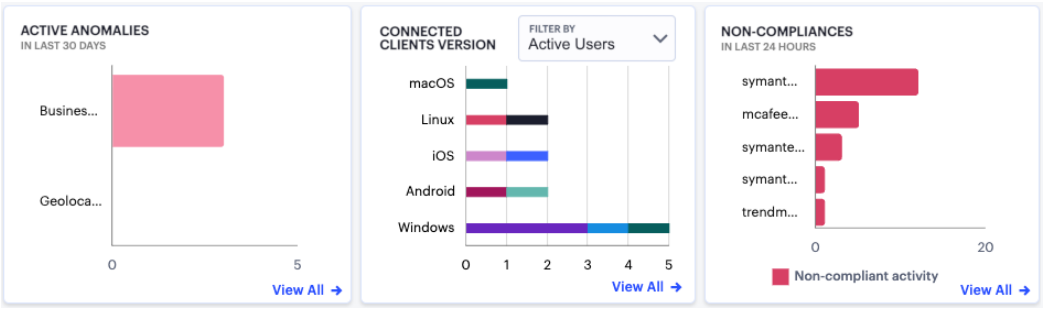

Viewing a breakdown of Active Anomalies, Connected Clients Versions, and Non-compliance

The **Active Anomalies** chart provides totals for the number of *Geolocation* anomalies and *Business Hours* anomalies. That is, application accesses that took place from an unexpected geographic location, or that took place outside of normal business hours. Hover your pointer over a particular bar to view a tooltip showing the label and total.

To view a detailed list of events that contributed to the totals in this chart, click **View all**:

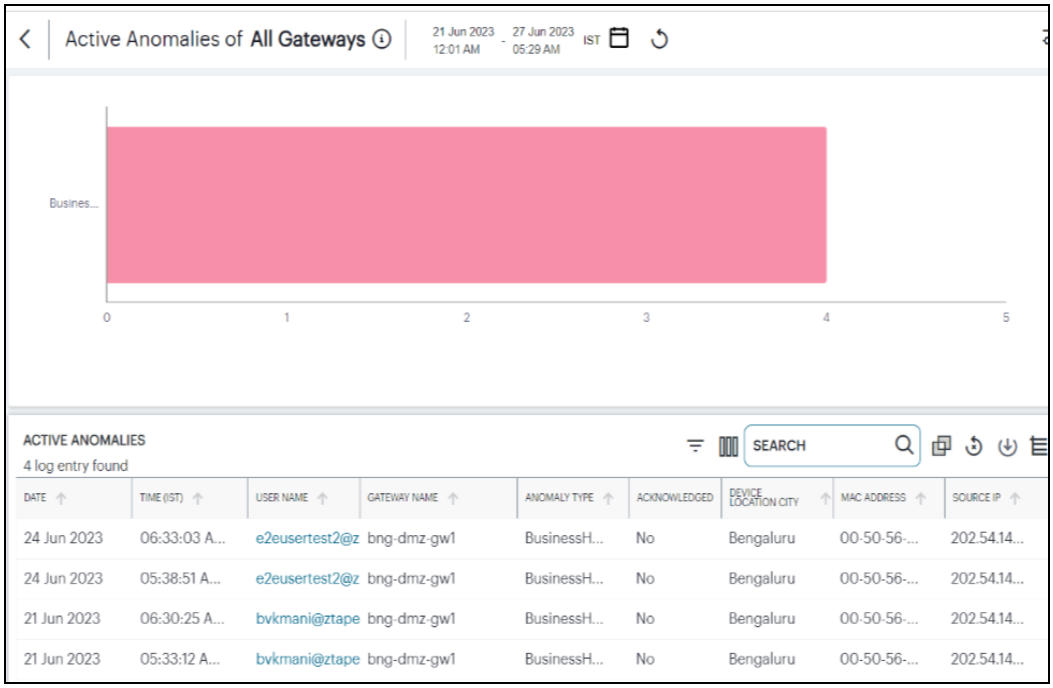

Viewing event logs for Active Anomalies

The **Connected Clients Version** chart shows totals for *Ivanti Secure Access Client* instances that have a current session on the *Controller*, broken down by device operating system type. Where more than one *Ivanti Secure Access Client* version is detected for a specific operating system, the bar is color-coded and relatively sized to represent each identified version and the number of clients using that version. Hover your pointer over a particular bar segment to view a tooltip showing the label and total.

If the currently selected time period is set to "Last Hour", this graph includes a drop-down control to filter the displayed data between:

- **Active Users**: Connected users during the last hour.
- **All Users**: Users that connected to the *Controller* in the last 30 days.

In all other time period views, the graph shows only data for all users connected during that time period.

To view a detailed list of events that contributed to the totals in this chart, click **View all**:

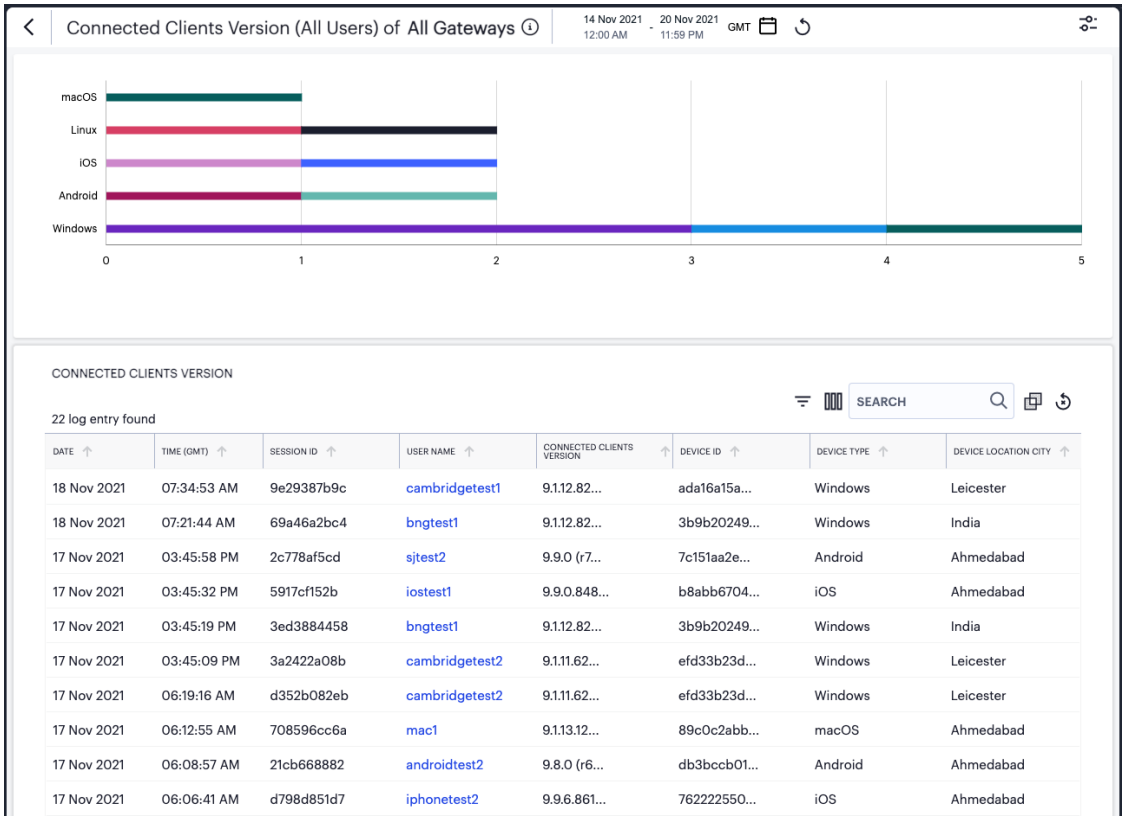

Viewing event logs for Connected Clients Versions

The **Non-compliance** chart provides a breakdown of non-compliant device activity that contravened a configured device policy. Totals are given for the highest policy contraventions recorded during the period.

To view a detailed list of events that contributed to the totals in this chart, click **View all**:

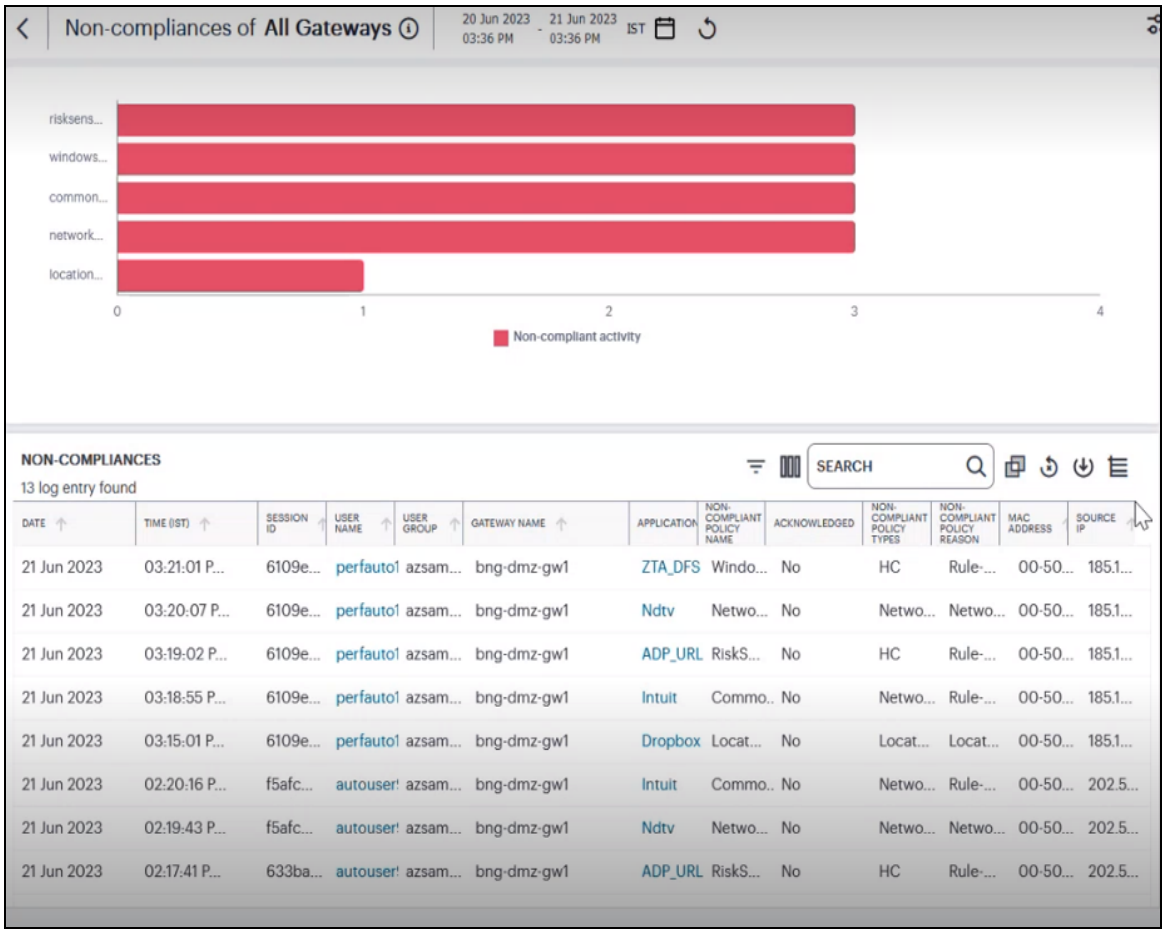

Viewing event logs for Non-compliances

To learn more about using the chart detail page, see Viewing Detailed Logs for a Chart.

Using the Top Active Breakdown Charts

The *radar* charts at the bottom of the page show a breakdown of **Gateways**, **User Locations**, and **Applications** across your organization. Each chart shows the *top active items* in each category.

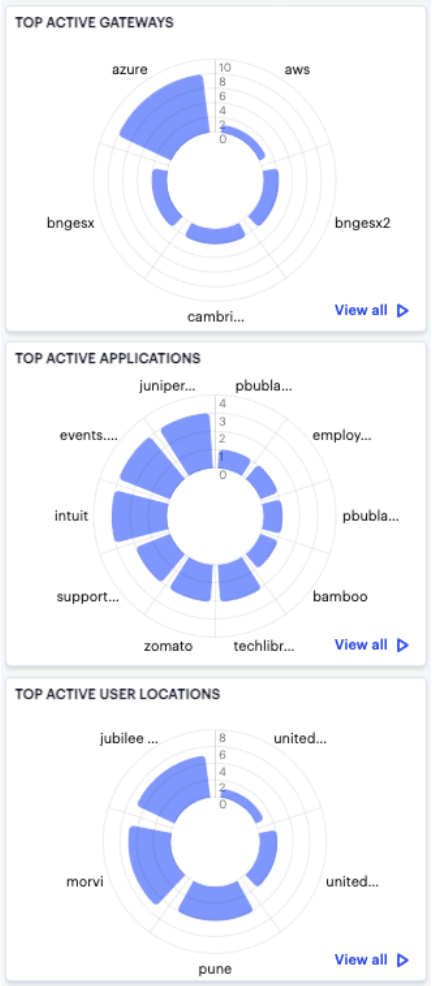

Viewing the Breakdown Radar Charts

Hover your pointer over a particular element to view a tooltip showing the label and total. To view more details and a set of log entries that constitute the date in the a chart, click the corresponding **View all** link. To learn more about using the chart detail page, see Viewing Detailed Logs for a Chart.

Viewing Detailed Logs for a Chart

When you select the **View all** link for any of the charts or graphs displayed within the *Insights* pages, the Tenant Admin portal displays a detail page containing a larger version of the selected chart, together with a table showing the event or log records that constitute the data points in the chart.

For example:

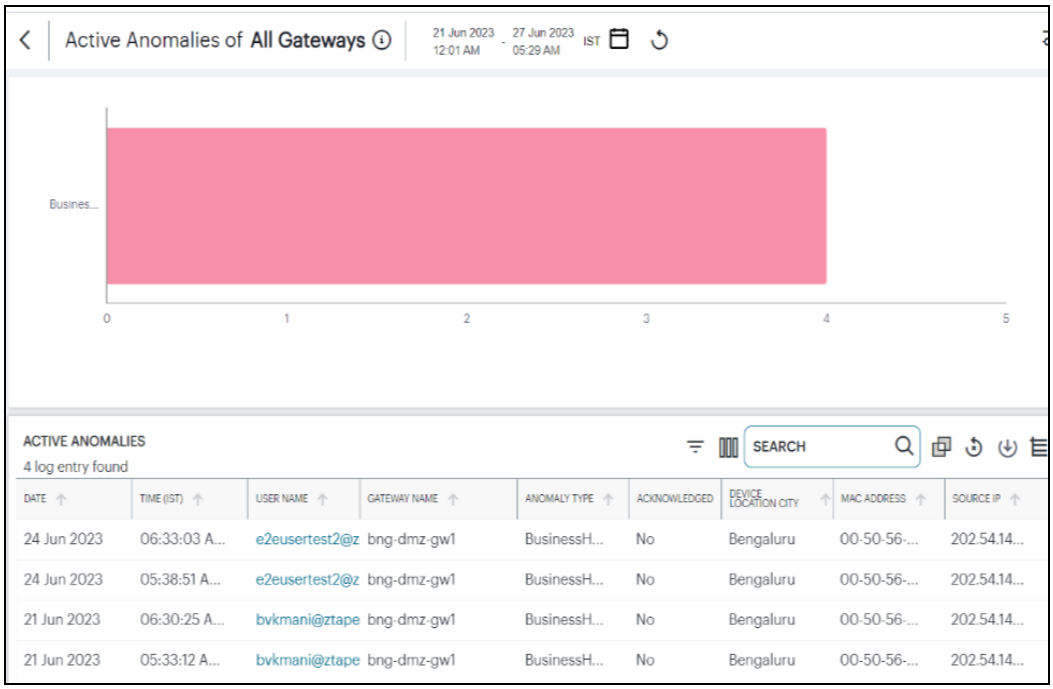

Viewing event logs for Active Anomalies

In this page:

- Hover your pointer over a specific bar in the chart to view a tooltip showing a numeric total for that category.
- Where a specific data item in the event table is truncated due to the column width, hover your pointer over the item to view a tooltip containing a full-length description. You can also re-size the width of any column by dragging the column.
- Select the name of a column to apply a sort by that criteria. The adjacent arrow shows either ascending or descending sort order, or no sort - select again to switch between each view.
- To copy a log's column content, double click on the column content and press **ctrl-c**. If the content has multiple words, for example log message column, then triple click and press **ctrl-c**. To paste the content elsewhere, press **ctrl-v**.

- To view a single log entry in a dedicated panel, click the corresponding three dot to activate the info-panel view. For example:

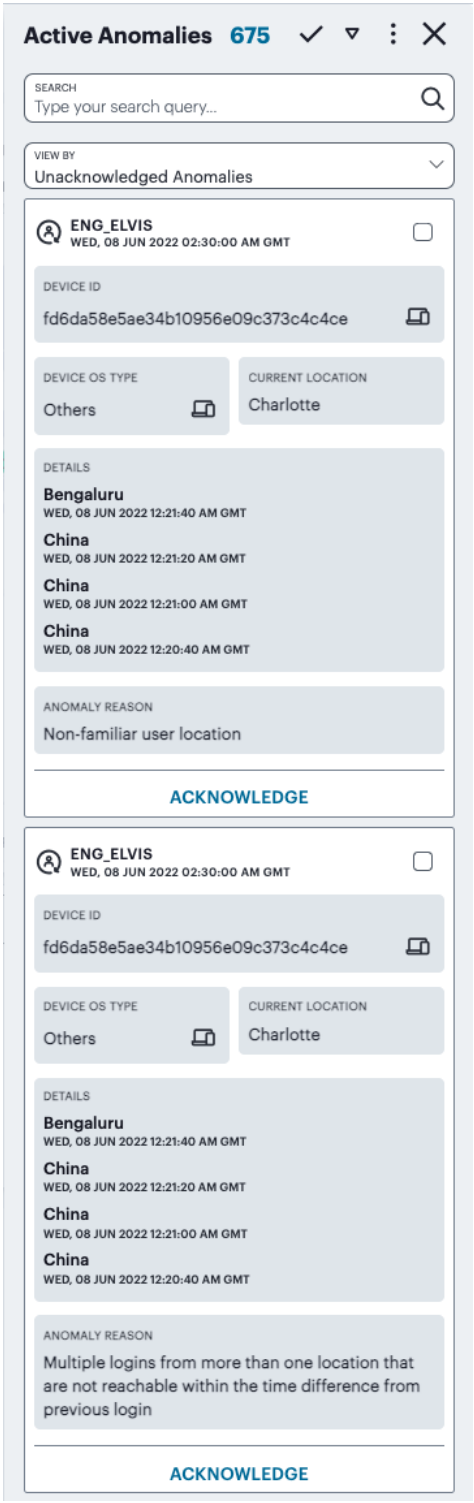

Viewing the Anomalies info panel

In the info-panel, use the **Previous** and **Next** icons to cycle through each event entry in turn.

- Use the date-time display at the top of the page (indicated) to apply a specific time period for the displayed data:

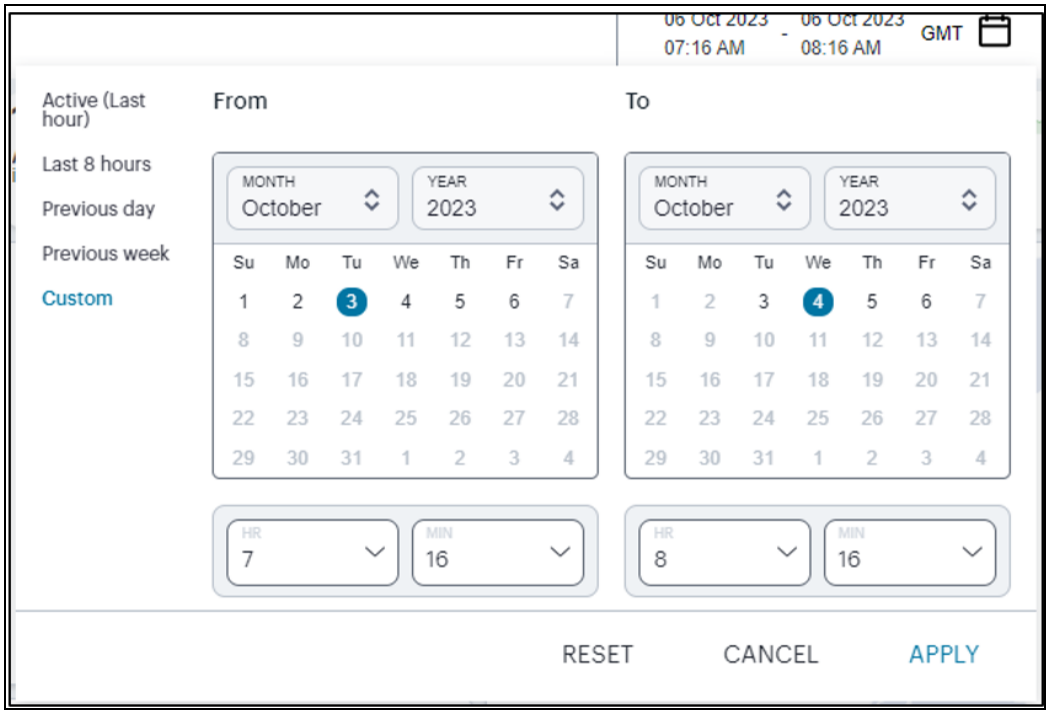

Selecting a date and time range

From the dialog, select the desired time period. Choose from the following predefined time period options:

- **Last hour**: Data observed for the previous 60 minutes.
- **Last <X> hours**: Data observed so far in the current day, up to the last hour (in GMT).
- **Previous day**: Data observed for the previous full day.
- **Previous Week**: Data observed for the previous calendar week (for the previous full Sunday-to-Saturday week).
- **Custom**: Data observed for a chosen time period. If you select this option, *nZTA* enables you to select a custom time period using the **From** and **To** date/time calendar controls.

The date/time calendar controls are enabled for only the **Custom** option. However, the calendar continues to identify the applicable start and end date-time for all predefined time periods.

To reset the selected time period back to the default (*Last Hour*) view, select **Reset**. To return to the current page without making any changes, select **Cancel**.

To apply your changes, select **Apply**. The selected time period is displayed in the filter bar and data across all **Insights** pages is updated accordingly.

- To manually refresh the display, select the following icon:

  

  Page refresh

- To search for a term in the displayed event data, select the following field:

  

  Search term highlighting

  nZTA highlights all matches in the event display.

- To trigger the advanced filter selection, use the following icon:

  

  Advanced Filtering

  To learn more, see Filtering the Logs.

- To export the displayed log as a CSV or JSON text file, or to create schedules to set up log export jobs, select the following icon:

  

  Export Filtered Logs

  To learn more, see Exporting Logs.

- To change the fields displayed for each event line, select the following icon:

  

  Show or hide event fields

  In the field selector, select a field name to toggle between show or hide. A *tick* icon indicates a displayed field. After you are finished, select the field selector icon to close the selector.

- To apply grouping to the event records, select the following icon:

  

  Group event records by selected criteria

This feature applies grouping to a selected field, such that event records are accumulated and grouped together under each unique data item identified in that field. Through grouping, an admin can quickly view the number of records of a particular type.

The criteria available for grouping depends on the chart being viewed, and reflects the field headings in that event table. For example, when viewing the **Top Active Applications** detail page (as shown above), you can choose to group by the following:

- Ungrouped
- User Name
- User Group
- Device Type
- Device ID
- Gateway Name
- Device Location City
- App Name

By selecting *App Name*, the event table is reconfigured to show a summary bar for each unique application identified in the logs.

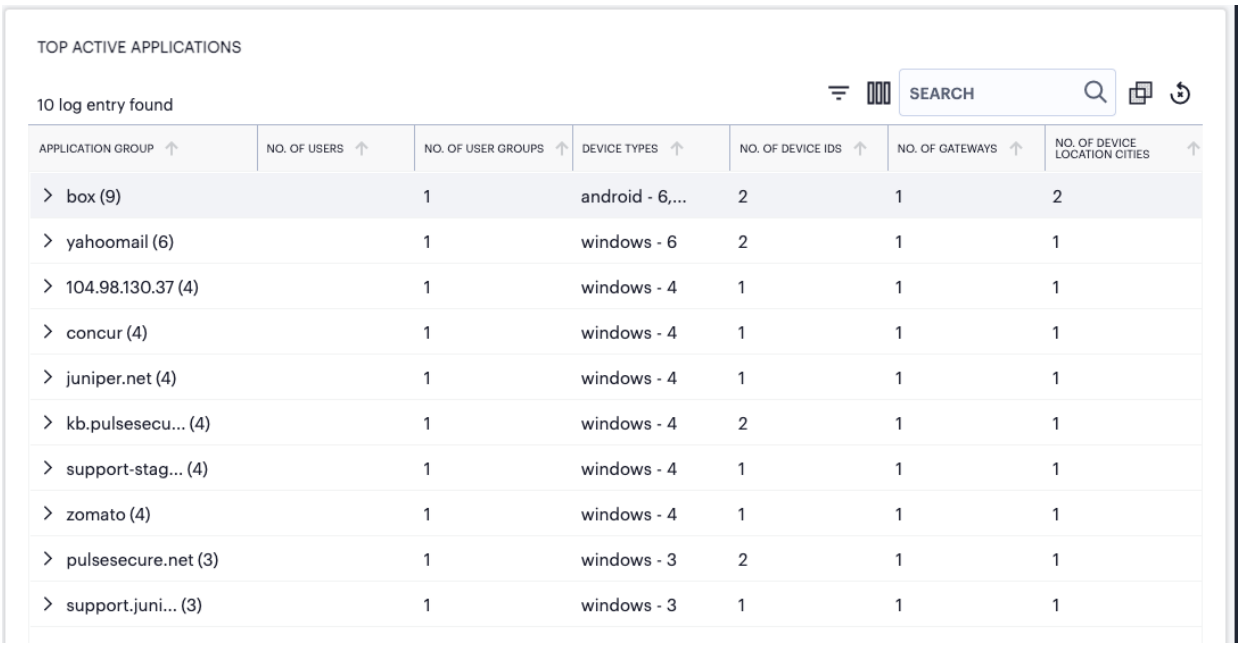

APPLICATION GROUP	NO. OF USERS	NO. OF USER GROUPS	DEVICE TYPES	NO. OF DEVICE IDS	NO. OF GATEWAYS	NO. OF DEVICE LOCATION CITIES
box (9)		1	android - 6,...	2	1	2
yahoomail (6)		1	windows - 6	2	1	1
104.98.130.37 (4)		1	windows - 4	1	1	1
concur (4)		1	windows - 4	1	1	1
juniper.net (4)		1	windows - 4	1	1	1
kb.pulsesecu... (4)		1	windows - 4	2	1	1
support-stag... (4)		1	windows - 4	1	1	1
zomato (4)		1	windows - 4	1	1	1
pulsesecure.net (3)		1	windows - 3	2	1	1
support.juni... (3)		1	windows - 3	1	1	1

Viewing Top Active Application events with grouping by application name applied

In this view, each application is identified together with a count of the number of event lines (in brackets) recorded against it. The event table field headers also adjust to reflect the summary counts recorded for each identified application.

To observe the event records in each grouping, select the arrow icon (indicated) adjacent to each application name in the table:

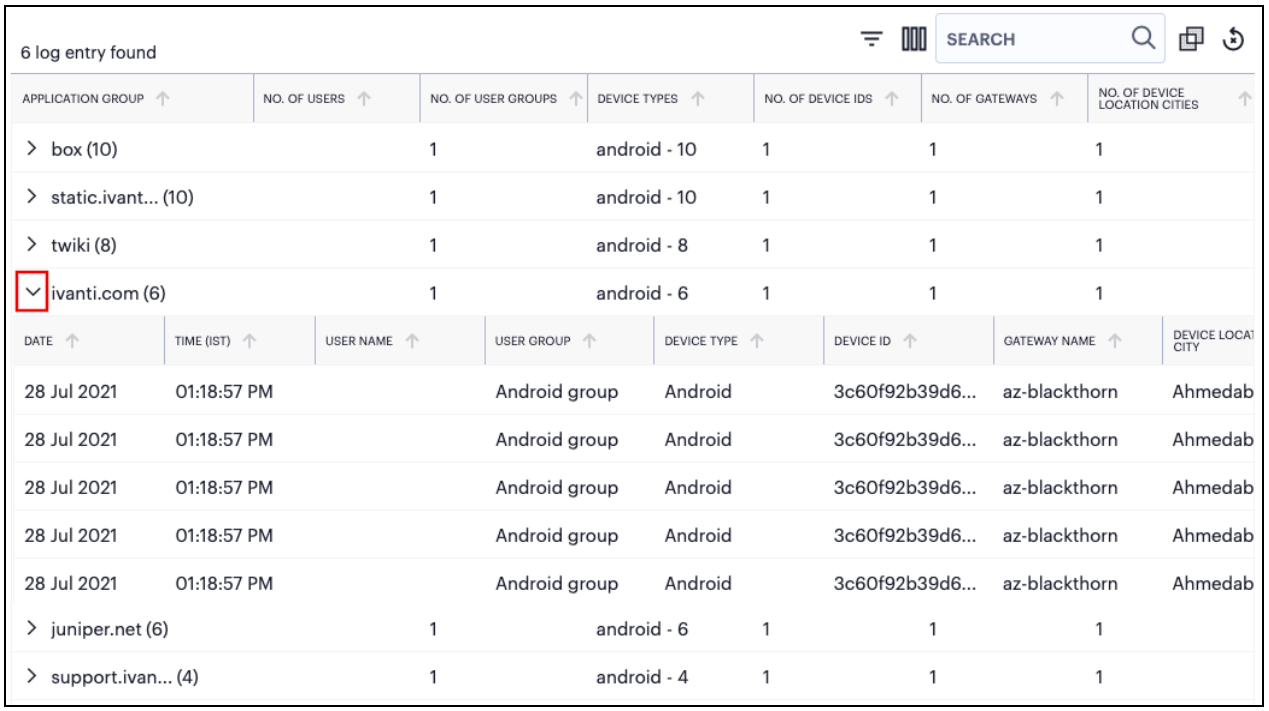

Viewing the grouped event records for a single named application

With grouping applied, the info-panel view adapts to reflect whether you selected a group header or an individual event record, such that the panel displays either the group totals or event details.

If you apply a grouping to the event data in this page, the page controls at the bottom (number of records per page and page navigation) have no effect.

- To remove any applied filters from the data set, select the following icon:

  Remove any applied filters from the data

- Use the page controls at the bottom of the window to select the number of event records/rows per page:

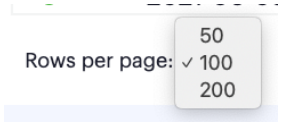

Setting the number of event rows per page

Choose from:

- 50

- 100 (default)
- 200

- To cycle through the event pages, use the page controls at the bottom-right.

Reviewing User Activity

User activity is available for all users, or for a specific user.

The *Users Overview* page shows activity relating to all users in your *nZTA* deployment.

To access the *Users Overview* page:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.

   The *Network Overview* page appears by default.

2. From the *nZTA* menu, click the **Insights** icon, then select **Users > All Users**.

The *Users Overview* page appears.

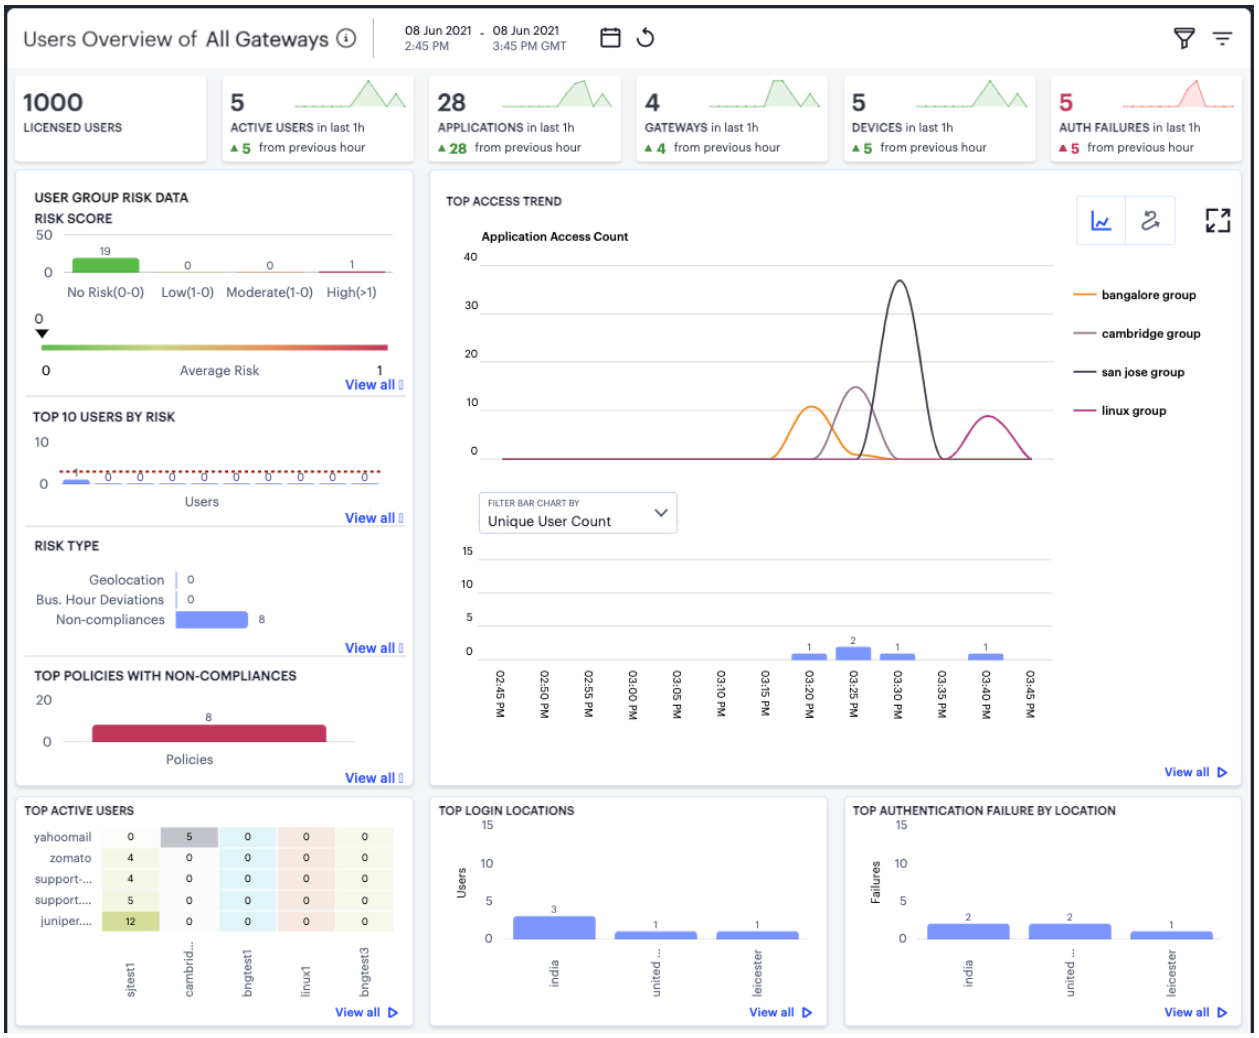

An overview of activity for all users

To view data relating to a specific user, see Showing Activity for a Specific User.

Understanding the Display

The *Users Overview* page contains the following components:

- **Filter bar**, allowing the selection of active or historic data. For details, see Using the Filter Bar.
- **Summary ribbon**, showing metrics for user activity. For more details, see Using the Summary Ribbon.
- **User Group UEBA Threat data**, showing graphs and metrics for UEBA Threat scores across your user groups. For more details, see Viewing a Summary of UEBA Threat Scores for your Users.

- **Top access trends**, showing a timeline chart of application access. For more details, see Viewing Top Access Trends.
- **Activity charts**, showing charts for *Top active users*, *Top login locations*, and *Top authentication failures by login location*. For more details, see Viewing User Activity Charts.

Each chart on this page includes a **View all** link. To learn more about using the chart detail page, see Viewing Detailed Logs for a Chart.

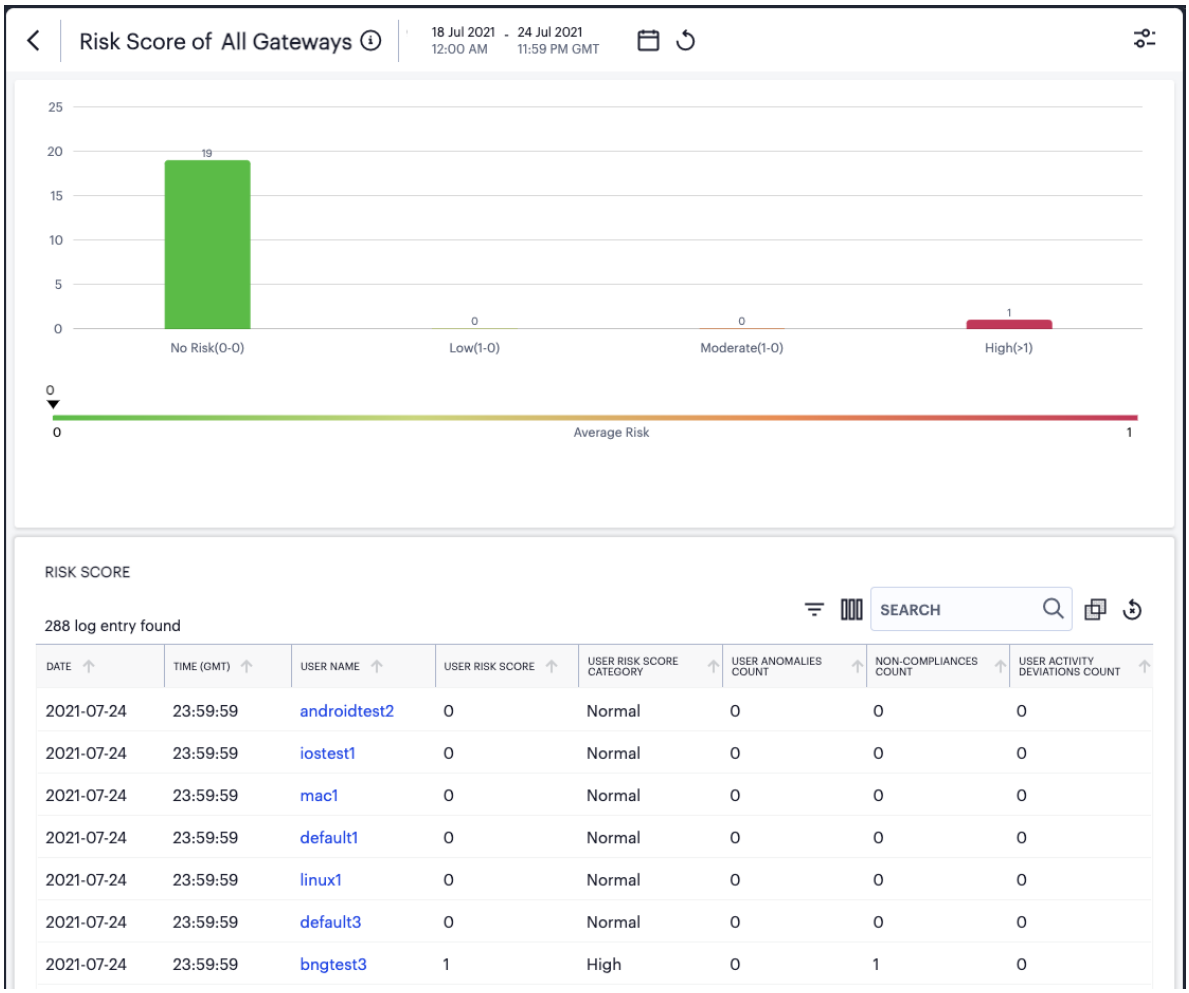

Viewing User Group risk detailed logs

Each detail view shows logs for the corresponding chart or category. To learn more about the detail page, including the features available, see Using the Active Anomaly, Connected Clients Version, and Non-Compliance Charts.

Using the Summary Ribbon

The Summary Ribbon at the top of the *Users Overview* page shows activity totals for the selected time filter:

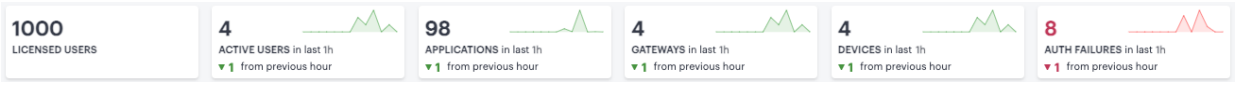

Viewing the summary ribbon

The ribbon indicates the totals accrued for each category during the displayed time period, as indicated adjacent to the category name. Hover your pointer over the category elements to show a descriptive tooltip.

- **Licensed Users**: The total number of licensed users.
- **Active Users**: The number of active users during the selected time period.
- **Applications**: The number of in-use applications.
- **Gateways**: The number of active *ZTA Gateways*.
- **Devices**: The number of active devices.
- **Auth failures**: The number of authentication failures.

By default, the data presented in the ribbon corresponds to the last hour. To change the time period, use the filter bar (see Using the Filter Bar).

If you are currently viewing data for the *last hour*, each category in the ribbon includes a trend graph (highlighted, top) showing the changes in data during the hour. Also included is a change value (highlighted, bottom) based on the previous hour:

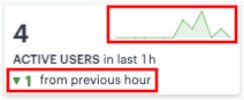

Data trends for last full hour versus the previous hour

If you select a historic time period in the filter bar, the ribbon displays only the main data totals for each category. Trend data is hidden.

Viewing a Summary of UEBA Threat Scores for your Users

On the **Insights** > **All Users** page, the *User Group UEBA Threat data* panel displays information concerning UEBA Threat factors across your user base:

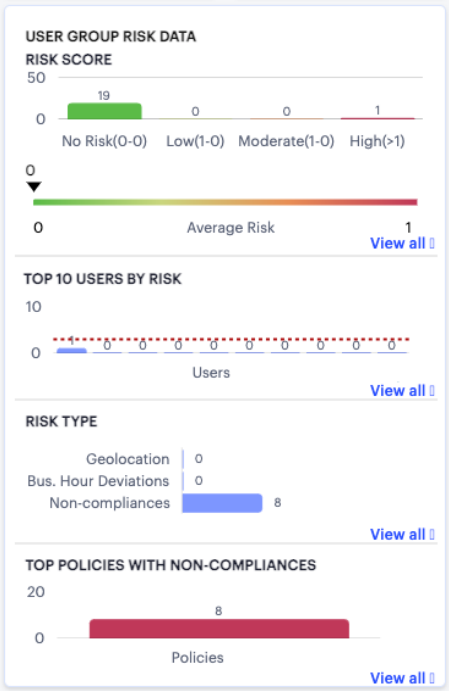

Viewing user group UEBA Threat data

The panel provides:

- A breakdown of UEBA Threat by user group.
- The average UEBA Threat score across all users.
- The top-10 users scoring highest for UEBA Threat.
- A break-down of UEBA Threat types.
- The policies with highest non-compliance.

A user's UEBA Threat score is calculated from a combination of three factors:

- Application access attempts originating from anomalous geographic locations or outside of normal business hours.
- Non-compliant user devices that attempted to access your applications.
- Activity Deviations.

Each additional incident increments a user's overall UEBA Threat score.

The *No. of users* chart provides a visual indication of the number of users that fall into each of the UEBA Threat categories. These categories are shown as percentage ratios of the overall UEBA Threat score and number of users. The upper and lower bands for each category are shown in brackets. The categories are:

- No risk (20%)
- Low (30%)
- Moderate (30%)
- High (20%)

> Where a particular UEBA Threat category matches no users for the selected time period, that category label is not shown.

Below this chart, *nZTA* displays the **Average UEBA Threat** score for all users on a scale between zero UEBA Threat and the highest UEBA Threat score measured at the end of the current time period.

> The maximum value shown in the chart corresponds to the highest UEBA Threat score for all users as they stand at the end of the time period, not the highest they have been within that period.

The *Top 10 Users by UEBA Threat* chart shows the top-10 users with the highest cumulative UEBA Threat score across the selected time period. Hover your pointer over each bar in the chart to see the name of the corresponding user. Where you have configured a UEBA Threat score action trigger (see Actions), this chart also contains a dotted line to indicate the UEBA Threat score threshold set in the action.

The *UEBA Threat Type* chart provides a breakdown of all geolocation anomalies, business hours deviations, and non-compliances that occurred during the selected time period.

The *Top Policies with Non-compliances* chart shows the device policies that recorded the highest number of non-compliances during the selected time period. Hover your pointer over each bar in the chart to see the name of the corresponding policy.

Viewing Top Access Trends

nZTA uses this section to show application access trends that occurred during the selected time period:

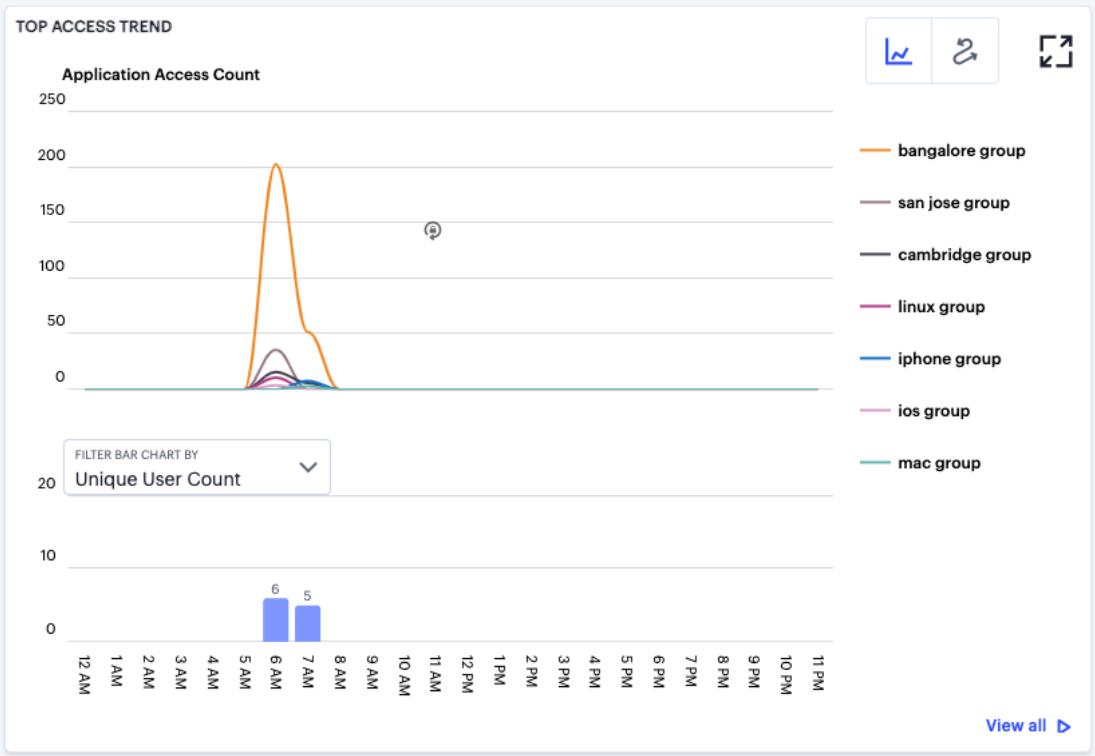

Viewing top access trends

You can choose to display this information through line and bar charts (as shown), or in an Sankey chart. Use the toggle icon at the top-right to select the required view:

Toggle between line/bar chart view and Sankey chart view

To expand the current view, click the Full Screen icon:

Expand the current view

In line/bar chart view, the display is split into two segments:

- A line chart showing the number of application accesses by each user group during each hourly period of the day
- A bar chart showing one of four data types, selected using the **Filter Bar Chart By** drop-down control:

- Unique User Count
- Unique Device Count
- Unique Location Count
- Anomalies

If you set a *Time Period* filter than spans more than one day, the data values shown in each hour period are cumulative totals for the same hour in each day during the time period.

In this chart, hover your pointer over each hourly interval to view a tooltip showing the corresponding data totals. Furthermore, you can click and drag a select box across a shorter time period to zoom in on a narrower time window. To return to the full 24 hour period, click the *zoom out* icon:

Zoom out from a selected time period

To toggle on or off the data for a particular user group, click the name of the group in the legend. Or, to view only the data for a specific user group, click the corresponding line in the graph.

In the Sankey chart view, *nZTA* provides an alternate visualization of application access activity, showing directed flow between related objects.

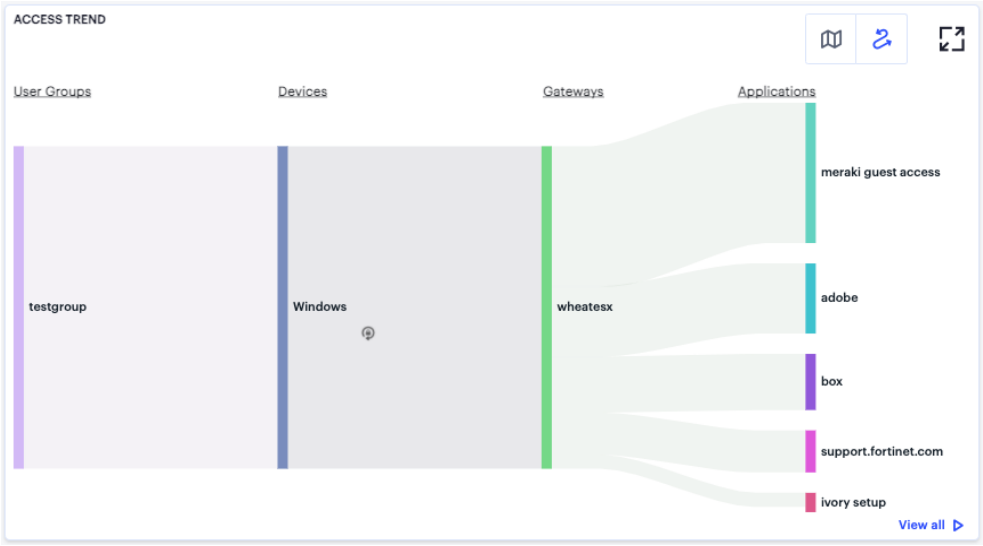

User Access Trends Sankey chart

The chart maps **User Groups** > **Devices** > **Gateways** > **Applications**. By hovering your pointer over a flow of interest, *nZTA* displays a tooltip confirming the scale of the activity between the two objects connected by the flow.

To focus the display on a specific flow, or to identify related objects that interact with this flow, click the chart at a point of interest. *nZTA* provides highlighting to all flows that pass through the point selected.

Viewing User Activity Charts

nZTA provides charts to represent user activity:

- **Top Active Users**: a grid showing users that accrued the highest number of successful accesses to your deployed applications. Tooltips show the number of accesses by a user for that application.
- **Top Login Locations**: a chart of the most active user locations per user group. Tooltips show a count of users active in that user group.
- **Top Authentication Failure by Location**: a chart showing totals for authentication failures observed per user location.
- **Top Risky Applications**: a chart showing the total users count for each of the top risky applications.

Hover your pointer over a particular element to view a tooltip showing the label and total.

Click **View All** to see the detailed metrics.

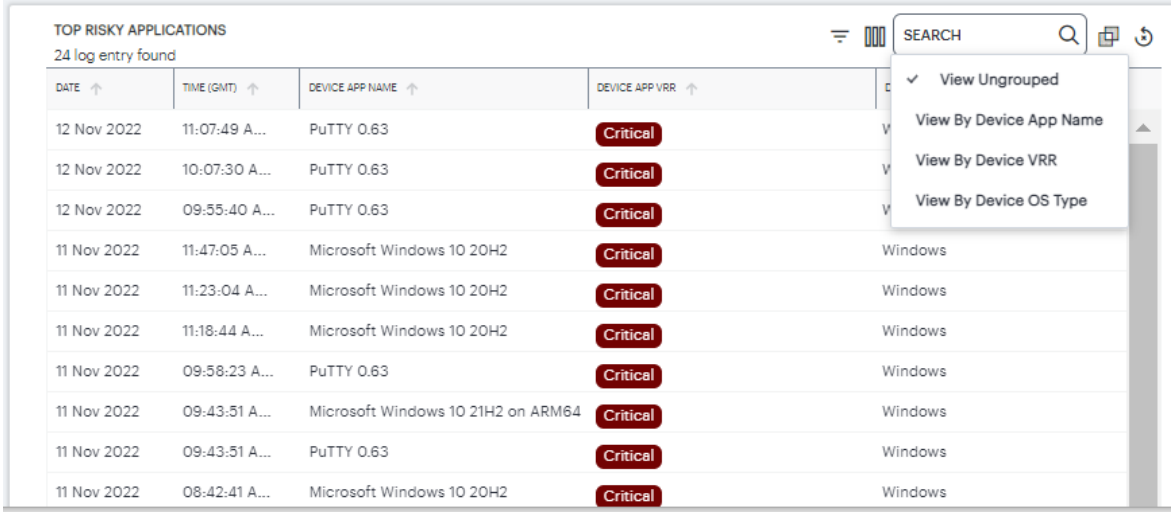

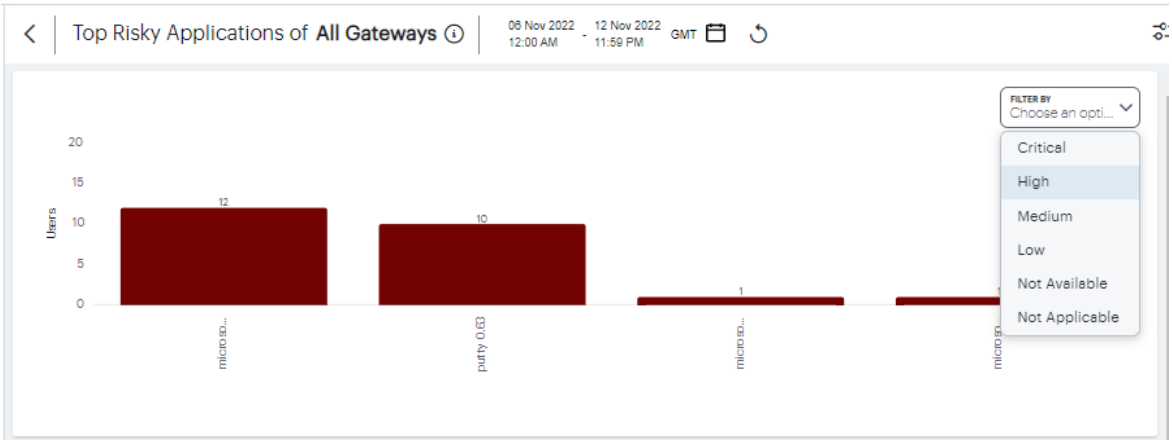

Top risky application details

Showing Activity for a Specific User

To view activity for a specific user, use the **Set User** option in the filter menu:

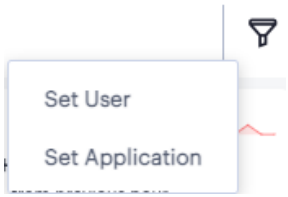

Activating the Set User option

Alternatively, from the **Network Overview** page, access specific user activity from the **Users** info-panel view. For more details, see Using the Summary Ribbon.

nZTA displays the *Users* page, showing activity for the selected user:

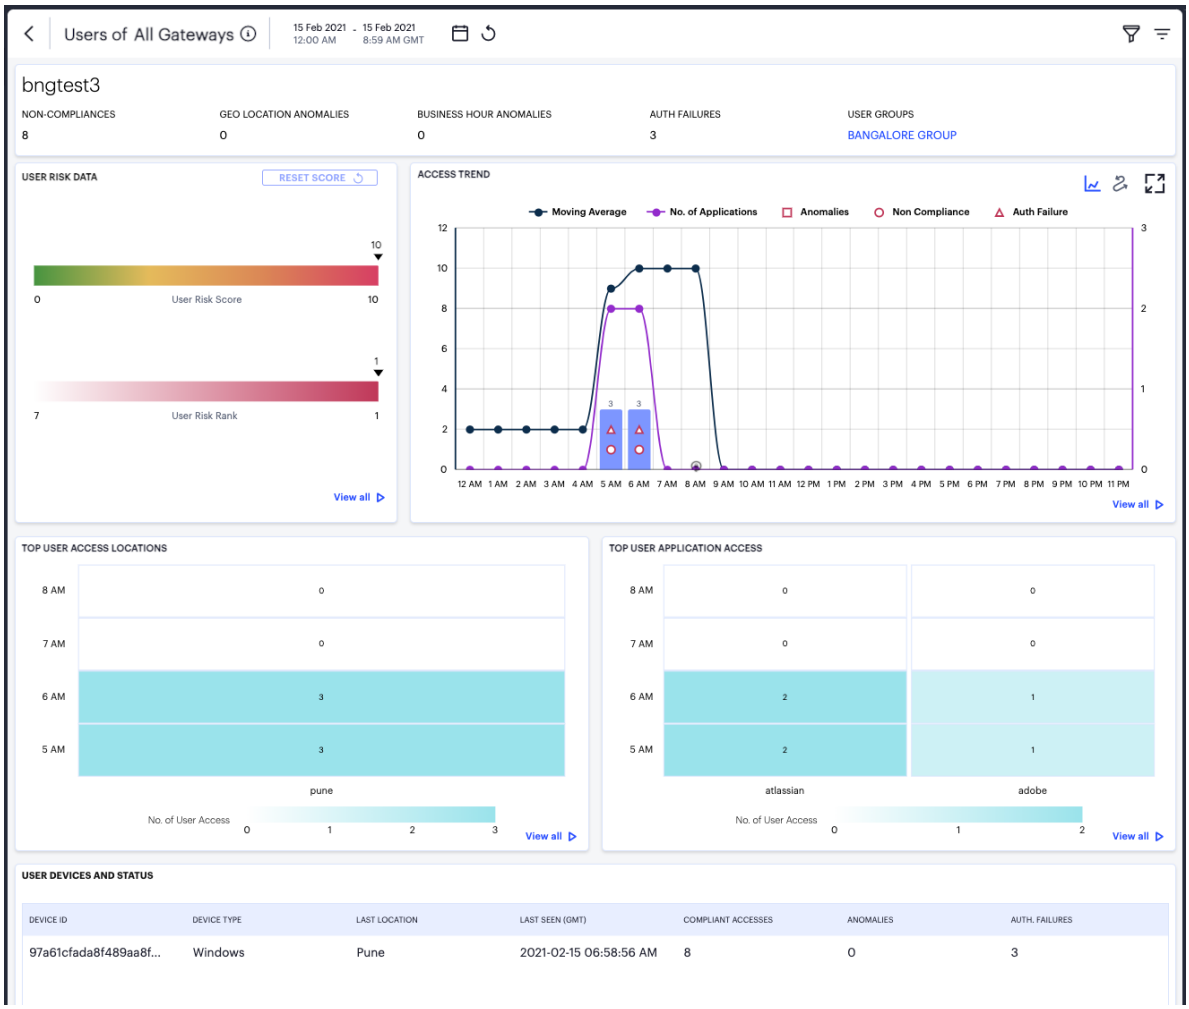

Viewing activity for a specific user

Understanding the Display

The *Users* page contains the following components:

- **Filter bar**, allowing the selection of active or historic data. For details, see Using the Filter Bar.
- **Summary ribbon**, showing activity metrics for the current user. For more details, see Using the Summary Ribbon.
- **User UEBA Threat data**, showing the User UEBA Threat Score and UEBA Threat Score Rank. For more details, see Viewing a Summary of UEBA Threat Scores for your Users
- **Access trend**, showing application accesses, non-compliance, and authentication failures by this

user over time.

- **Activity charts**, showing top user access locations and application activity.

Each chart on this page includes a **View all** link. This link provides access to a detail view showing logs for the corresponding chart. For example:

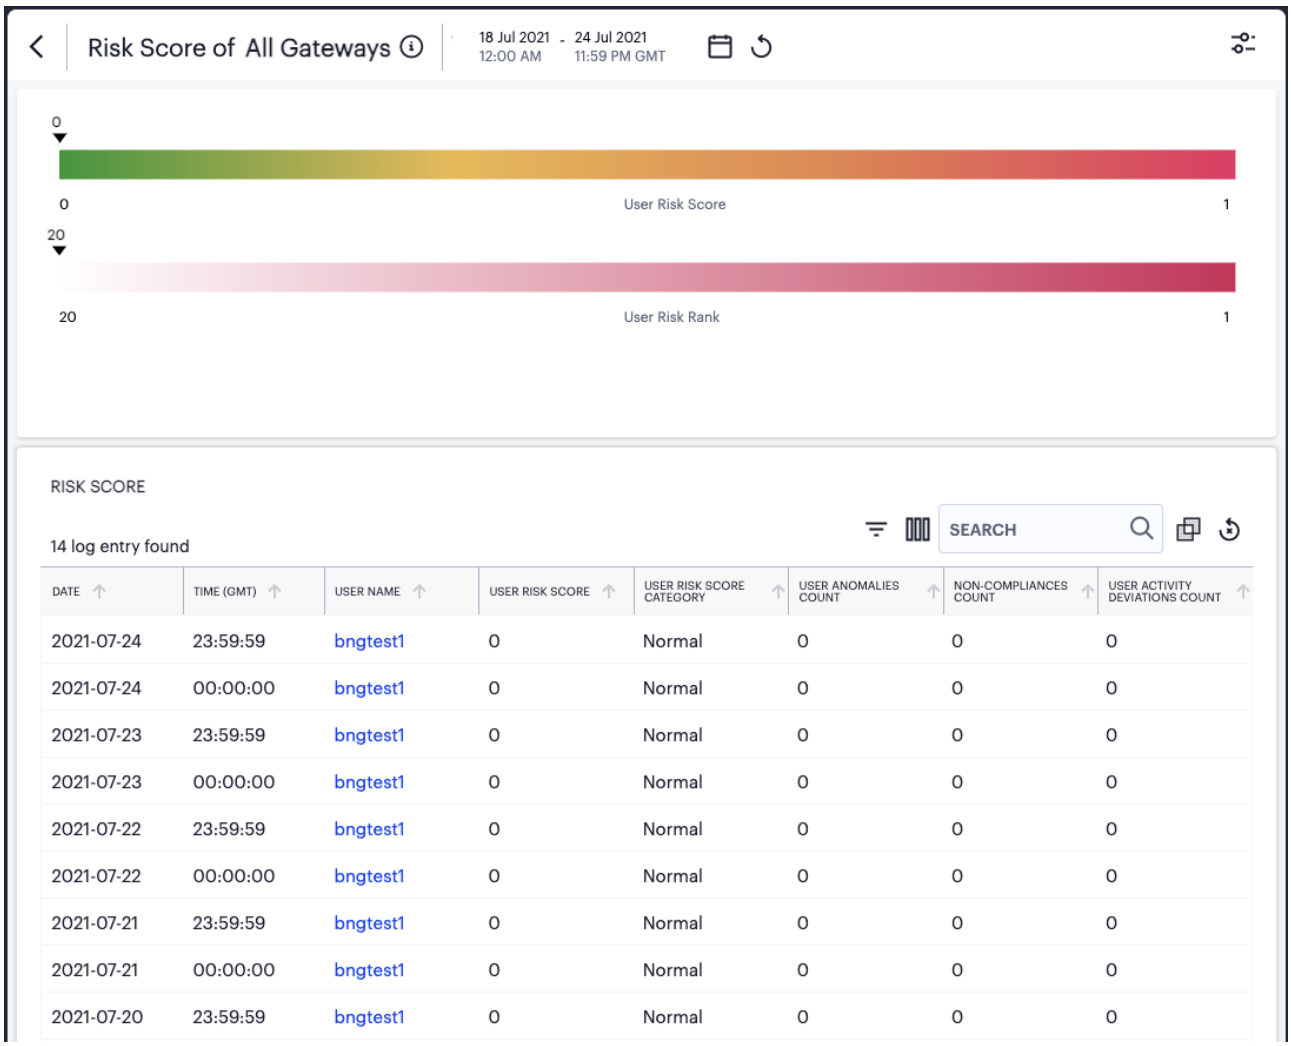

Viewing User UEBA Threat Score detailed logs

Each detail view shows logs for the corresponding chart or category. To learn more about using the chart detail page, see Viewing Detailed Logs for a Chart.

Using the Summary Ribbon

The Summary Ribbon at the top of the *Users* page shows activity totals for the user during the selected time filter:

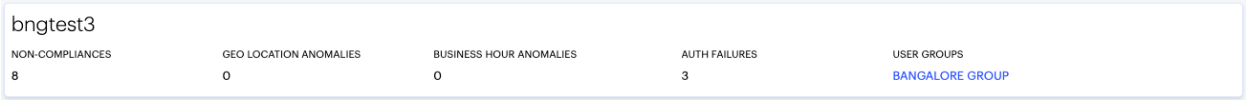

Viewing the summary ribbon

The ribbon indicates totals accrued for the selected user during the displayed time period. The summary ribbon provides the following metrics:

- **Non-compliances**: The number of non-compliant access attempts by this user during the period.
- **Geo Location Anomalies**: The number of application accesses attempted from anomalous geographic login locations by this user during the period.
- **Business Hours Anomalies**: The number of application accesses attempted outside of normal business hours by this user during the period.
- **Auth failures**: The number of authentication failures experienced by this user during the period.
- **User groups**: The user groups of which this user is a member. Click the name of a group to access the *user groups* page.

Viewing UEBA Threat Data for the Selected User

The *User UEBA Threat Data* panel displays information concerning UEBA Threat for the selected user:

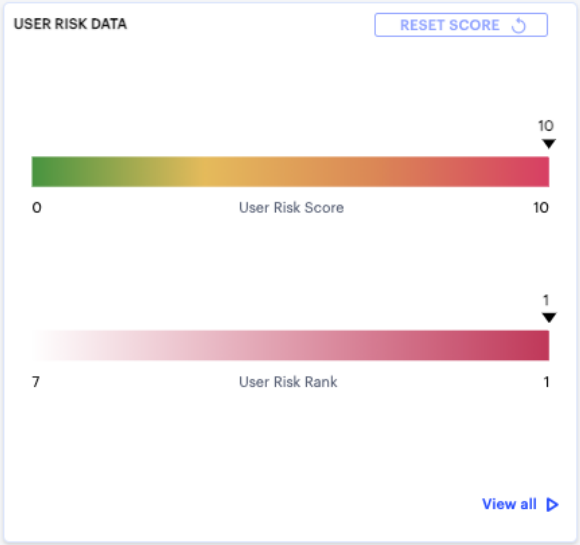

Viewing UEBA Threat data for a user

The panel provides:

- The selected user's *UEBA Threat score*, as calculated at the end of the selected time period.

  The UEBA Threat score is shown as an indicator on a linear scale of no risk up to the highest recorded score during the time period. To learn more about a user's UEBA Threat score, see Viewing a Summary of UEBA Threat Scores for your Users.

- The selected user's *UEBA Threat Score rank*, as calculated at the end of the selected time period.

  A user's UEBA Threat Score rank is the UEBA Threat score as measured against other active users in the organization, displayed on a linear scale. As a user increases their UEBA Threat score, the more the rank position (the indicator) decreases towards 1 out of the total of active users (the value at the start of the scale). A rank of "1" means that a user ranks highest for risk out of all active users.

- A link to reset the selected user's UEBA Threat score

Viewing Access Trends for the Selected User

nZTA uses this section to show access trends for the selected user that occurred during the selected time period.

You can choose to display this information through line and bar charts, or in a Sankey chart. Use the toggle icon at the top-right to select the required view:

Toggle between line/bar chart view and Sankey chart view

To expand the current view, click the Full Screen icon:

Expand the current view

The line and bar chart shows user access trends through each hourly interval of the day:

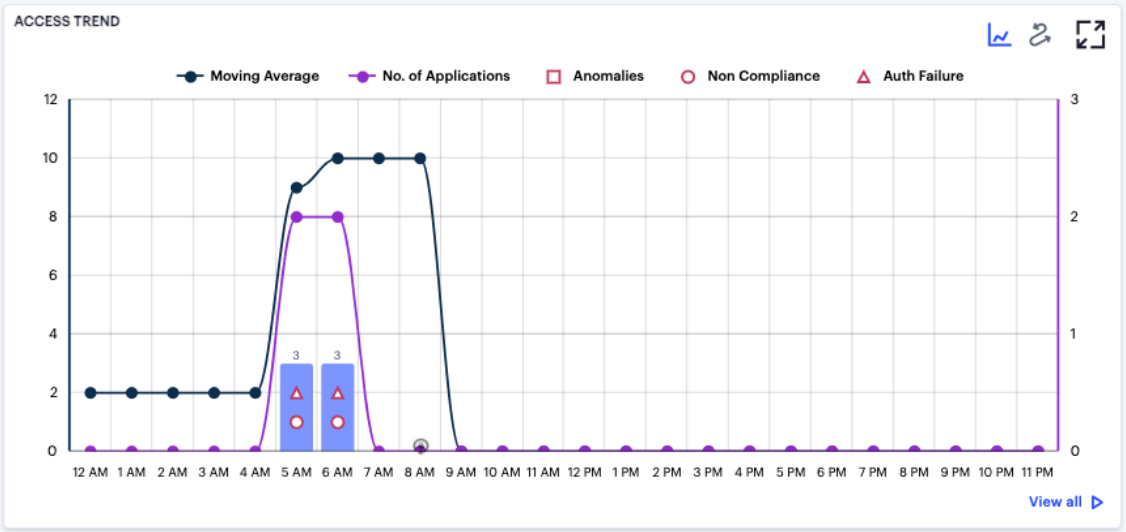

Viewing hourly access trends for a user

The horizontal axis reflects the 24 hourly intervals in a day, with the vertical axis showing the number of accesses. As a user interacts with your *nZTA* infrastructure, access attempts are recorded and shown in this chart according to the hour in which they occurred. This in turn provides an overview of the daily access trends for the user.

The bars denote the number of accesses made, with the lines charting the number of applications accessed, and the moving average over 30 days.

Anomalies, non-compliances, and authentication failures detected in the hour are marked as per the legend.

Hover your pointer over an hour interval to see a tooltip summary of data points for that hour:

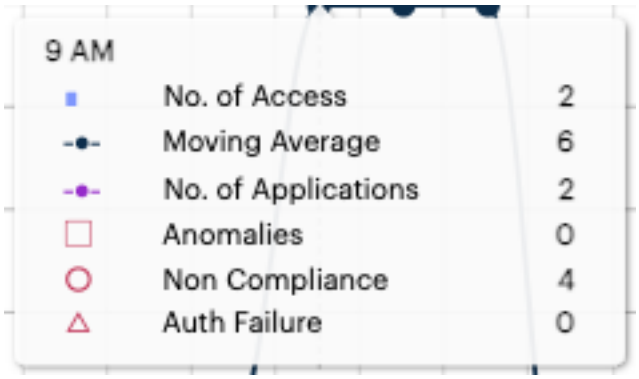

Viewing a data summary tooltip

To learn more about the User Access Sankey chart, see Viewing the User Access Sankey Chart.

Viewing the User Access Sankey Chart

The User Access Sankey chart provides an alternate visualization of your selected user's activity, showing directed flow between related objects. The width of each stream in the flow is proportional to the utilization of the object the flow passes through, allowing an administrator to view significant usage trends for the selected user and your application infrastructure.

To toggle between the User UEBA Threat Score chart and the User Access Sankey chart, use the icons at the top-right:

Toggle between User UEBA Threat Score view and User Access Sankey chart view

By clicking the toggle display icon, the User Access Sankey chart replaces the User UEBA Threat Score graph in the display. All other components remain unchanged.

Displaying the User Activity Sankey Chart View

The *nZTA* User Activity Sankey chart maps **Locations** > **Devices** > **Gateways** > **Applications** for the selected user. By hovering your pointer over a flow of interest, *nZTA* displays a tooltip confirming the scale of the activity between the two objects connected by the flow.

To focus the display on a specific flow, or to identify related objects that interact with this flow, click the chart at a point of interest. *nZTA* provides highlighting to all flows that pass through the point selected.

To activate the User Non-Compliances Sankey chart view, use the toggle icons at the bottom of the chart:

Toggle between User Access Sankey chart view and the User Non-Compliances Sankey chart view

Use this toggle to switch the Sankey chart between displaying User Application Access or User Non-Compliances flows.

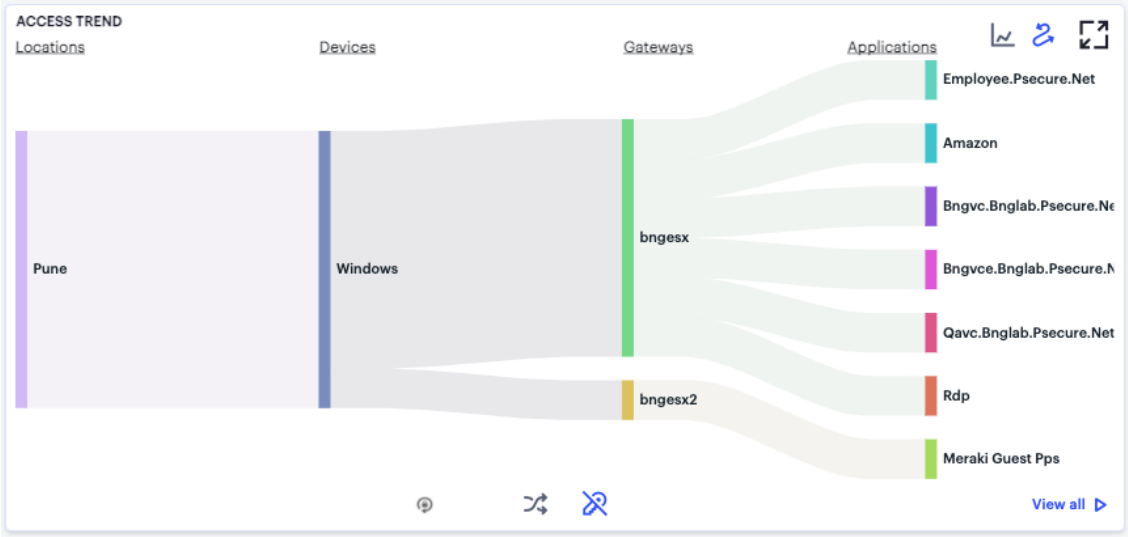

Displaying the User Non-Compliance Sankey Chart View

Viewing User Activity Charts

The **Top User Locations** and **Top User Activity** charts show the top locations and applications the user is active with at different times of the day. Each chart provides a visual breakdown of normal activity across the day, with anomalies highlighted when they occur.

Viewing and Terminating User Sessions

To view the list of currently active user sessions:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.

   The *Network Overview* page appears.

2. From the *nZTA* menu, click the **Insights** icon, then select **Users > User Sessions**.

   Viewing active user sessions

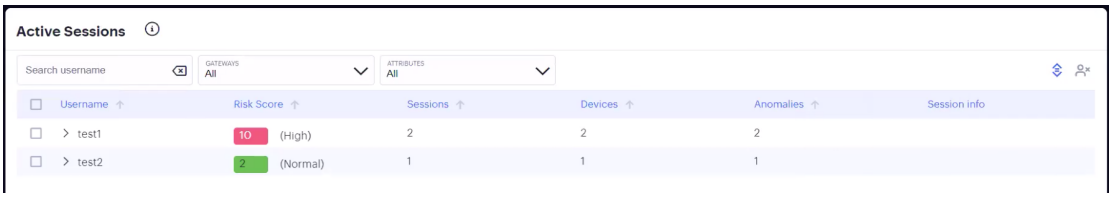

   The *User Sessions* page appears:

Use this page to view currently-active user sessions, and to terminate selected sessions as required. Each row corresponds to a single user and shows the following details:

- The username
- The user's UEBA Threat score. For more information on UEBA Threat scores, see Reviewing User Activity.
- The number of active sessions.
- The number of devices used.
- The number of anomalies observed. For more information on anomalies, see "Using the Active Anomaly, Connected Clients Version, and Non-Compliance Charts" on page 613.
- Session information, if available

Click the arrow icon adjacent to each column to sort in ascending or descending order.

Use the search boxes at the top of the page to search by:

- an entered **username**
- a specified **Gateway**
- **attributes**:
  - *Username*: enter a user name
  - *Device ID*: select a device ID
  - *Risk*: select a UEBA Threat score level

The data automatically updates to reflect the chosen search criteria.

Click the arrow icon adjacent to the user name to view all active sessions for the user:

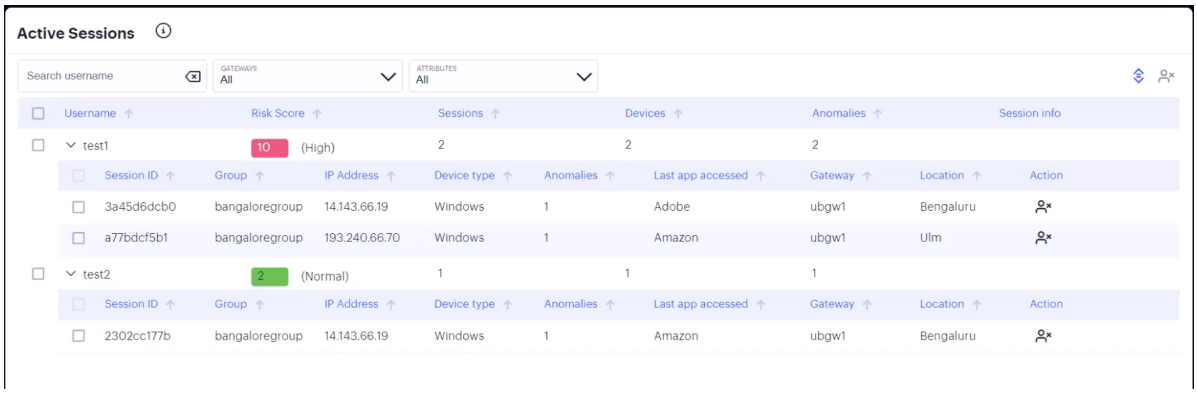

Viewing all active sessions for a user

Alternatively, to expand or collapse the list of sessions for *all* users, click the icon at the top-right:

Expand or collapse the complete user session list

To terminate a specific user session, locate the session row on the page and click the corresponding *terminate* button:

Terminate a session

To terminate multiple sessions in one operation, use the checkboxes adjacent to each username (to terminate all sessions for that user), or adjacent to each session row (to terminate individual sessions for one or more users). Then click the *terminate multiple sessions* button at the top-right.

In all cases, *nZTA* provides a confirmation dialog showing the session, or sessions, selected to be deleted:

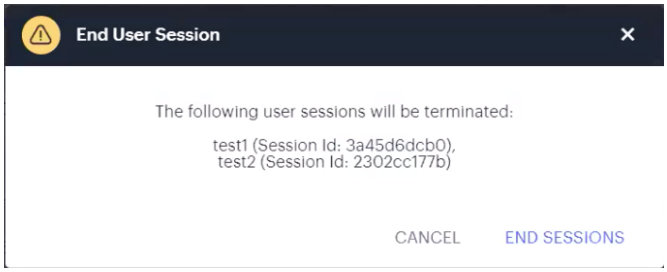

Confirming the session(s) to terminate

All session terminations performed through this page are logged in the *nZTA* Access Logs.

You can also terminate active user sessions through the Anomalies info-panel. For more details, see Using the Summary Ribbon.

Reviewing Application Usage

Applications in *nZTA* are defined primarily by the URI you use to access them, and can be fully *defined* (for example, a complete URI denoting a specific application at a location) or *discovered* (for example, a wildcard-prefixed FQDN that denotes an endpoint containing one or more applications).

The *Insights* > *Applications* pages shows usage data for all applications requested through your *nZTA* deployment.

nZTA provides the following views for your application usage:

- **All Applications**: Shows usage metrics for all defined applications in your *nZTA* deployment. See All Applications.
- **Discovered Applications**: Shows usage metrics for all discovered applications in your *nZTA* deployment. See Discovered Applications.
- **Default Gateway Applications**: Shows usage metrics for all applications managed through the default *nZTA Gateway* defined in your *Application Discovery* secure access policy. See Default nZTA Gateway Applications.

> A default *nZTA Gateway* is used to handle all requests from applications that are not referenced by any secure access policy. To learn more about setting a default *nZTA Gateway*, see Configuring a Default Gateway for Application Discovery.

To learn more about defining applications for use with secure access policies, see Defining Applications and Application Groups.

To view application usage:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.

   The *Network Overview* page appears.

2. From the *nZTA* menu, click the **Insights** icon, then select **Applications** and choose either **All Applications**, **Discovered Applications**, or **Default Gateway Applications**.

   The *Applications Overview* page appears, showing the selected metrics.

Use the tabs at the top of the page to switch between the different views: *All Applications*, *Discovered Applications*, and *Default Gateway Applications*. Each tab consists of a number of panels containing metrics and charts to show application usage in one of the aforementioned categories.

To view data relating to a specific application, see Showing Usage Data for a Specific Application.

All Applications

The *All Applications* tab shows usage metrics for all defined and discovered applications:

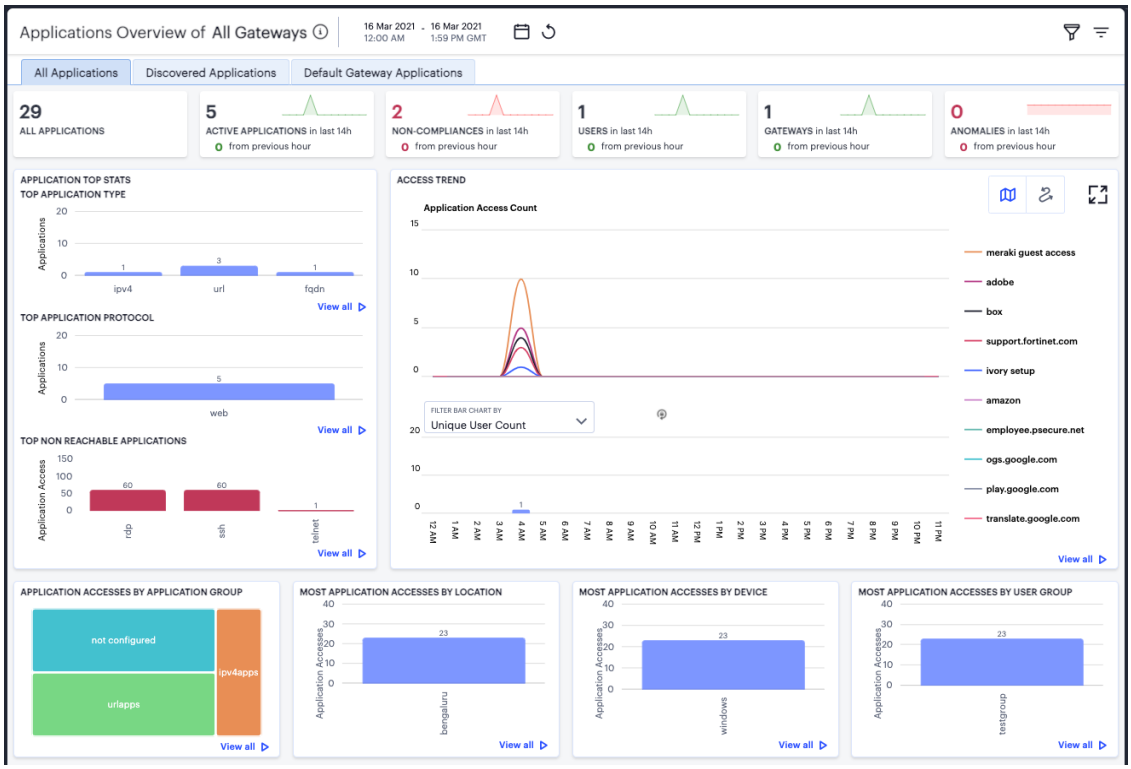

Viewing usage charts and graphs for your applications

The display is split into sections:

- **Summary Ribbon**
- **Application Top Stats**
- **Access Trends**
- **Activity charts** for *Application Accesses by Application Group*, *Most Application Accesses by Location*, *Most Application Accesses by Device*, and *Most Application Accesses by User Group*.

Each chart in the display includes a **View all** link providing access to a detail page showing log records for the corresponding chart. These log records include links to the application and user involved in the logged event. *Ivanti* recommends using this page to access the metrics page for the specific application (see Showing Usage Data for a Specific Application) or user (see Showing Activity for a Specific User). This method of navigation offers an alternative to searching for a specific application through the "Select Application" filter option, where the exact application name might not be known (for example, discovered and default applications not specifically defined in a secure access policy). To learn more about using the chart detail page, see Viewing Detailed Logs for a Chart.

The **summary ribbon** provides the following metrics:

- **All Applications**: The total number of applications defined on the *Controller*.
- **Active Applications**: The number of applications accessed during the selected time period.
- **Non-compliances**: The number of non-compliant attempts to access applications.
- **Users**: The number of active users.
- **Gateways**: The number of active *nZTA Gateways*.
- **Anomalies**: The number of anomalous application accesses based on geographic and business hours irregularities.

The **Application Top Stats** panel provides the following charts:

- **Top Application Type**: A bar chart showing the application types that attracted the greatest numbers of application accesses during the selected time period (for example, FQDN, URL, or IP address).
- **Top Application Protocol**: A bar chart showing the application protocol types that attracted the greatest number of application accesses during the selected time period (for example, Web, RDP, or SSH).
- **Top Non Reachable Applications**: A bar chart showing the applications marked most-often not reachable by the *Controller*. To learn more about application availability status, see Viewing your Secure Access Policies.

The Top Non Reachable Applications chart includes only applications where the status can be determined. It does not show applications where the status is unknown, such as for applications based on FQDNs, wildcard-based FQDNs, and IPv4/IPv6 ranges - all of which are unsupported by the application health monitoring feature.

For all charts, hover your pointer over each bar to display a tooltip of the type and number of accesses recorded.

The **Access Trends** panel shows application access trends that occurred during the selected time period. You can choose to display this information through line and bar charts, or in a Sankey chart. Use the toggle icon at the top-right to select the required view:

Toggle between line/bar chart view and Sankey chart view

To expand the current view, click the Full Screen icon:

Expand the current view

Click this icon again to return to the standard view.

In line/bar chart view. The display is split into two segments:

- A line chart showing the number of accesses for the top-10 applications during each hourly period of the day
- A bar chart showing one of four data types, selected using the **Filter Bar Chart By** drop-down control:
- Unique User Count: Shows a count of unique user activity identified during each hourly period.
- Unique Device Type Count: Shows a count of unique device types identified during each hourly period.
- Unique Location Count: Shows a count of activity from unique user locations identified during each hourly period.

- Unique User Group Count: Shows a count of activity from unique user groups identified during each hourly period.

> If you set a Time Period filter than spans more than one day, the data values shown in each hour period are cumulative totals for the same hour in each day during the time period.

In this chart, hover your pointer over each hourly interval to view a tooltip showing the corresponding data totals. Furthermore, you can click and drag a select box across a shorter time period to zoom in on a narrower time window. To return to the full 24 hour period, click the zoom out icon:

Zoom out from a selected time period

To toggle on or off the data for a particular application, click the name in the legend. Or, to view only the data for a specific application, click the corresponding line in the graph.

In the Sankey chart view, *nZTA* provides an alternate visualization of application access activity, showing directed flow between related objects.

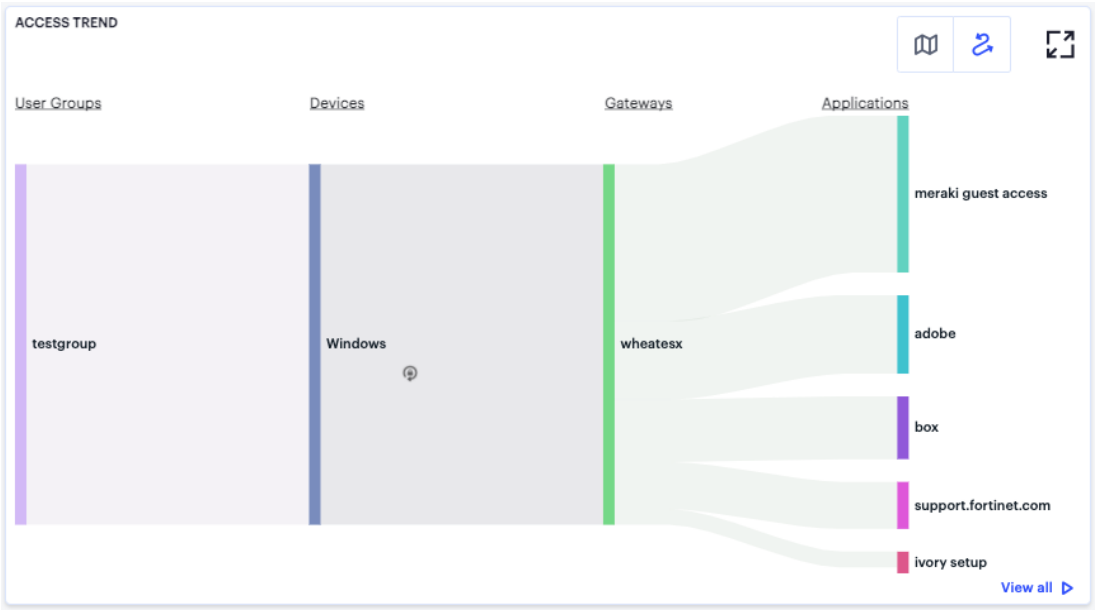

User Access Trends Sankey chart

The chart maps **User Groups > Devices > Gateways > Applications**. By hovering your pointer over a flow of interest, *nZTA* displays a tooltip confirming the scale of the activity between the two objects connected by the flow. To focus the display on a specific flow, or to identify related objects that interact with this flow, click the chart at a point of interest. *nZTA* provides highlighting to all flows that pass through the point selected.

The **Activity Charts** on this page represent top application access totals in the following categories:

- **Application Accesses by Application Group**: a grid chart showing the application groups containing the applications that accrued the highest number of successful accesses. Application group sizes in the chart are proportional to the number of accesses, compared with other groups. Tooltips show a count of the accesses made to that group. To learn more about Application Groups, see Adding Application Groups to the Controller.
- **Most Application Accesses by Location**: a bar chart showing a list of the most active user locations with respect to application access. Tooltips show a count of the application accesses by that location.
- **Most Application Accesses by Device**: a bar chart showing a list of the most active user device types with respect to application access. Tooltips show a count of the application accesses by that device type.
- **Most Application Accesses by User Group**: a bar chart showing a list of the most active user groups with respect to application access. Tooltips show a count of the application accesses by users in that user group.

Hover your pointer over a particular element to view a tooltip showing the label and total.

Discovered Applications

The *Discovered Applications* tab shows usage metrics for applications discovered by the *Controller* for applications defined with a wildcard domain and with **Application Discovery** enabled:

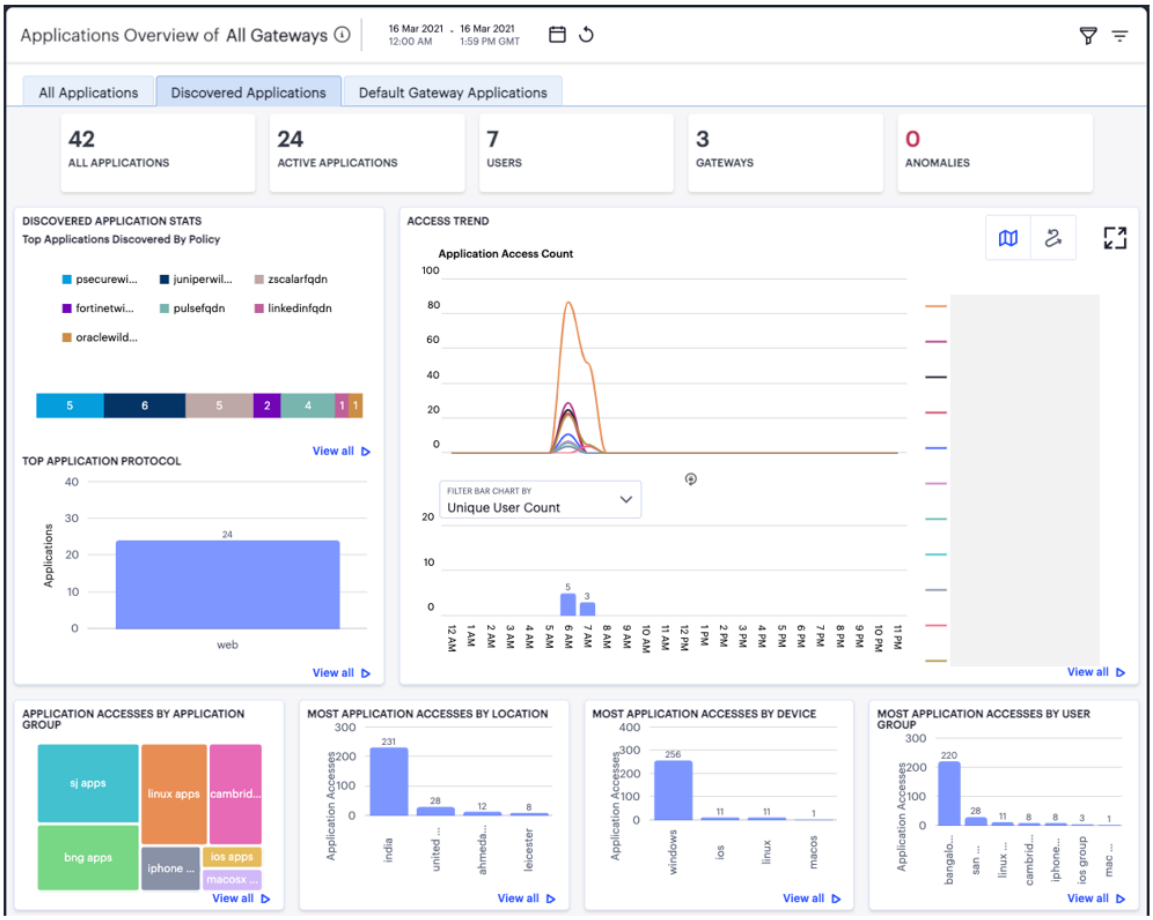

Viewing usage charts and graphs for discovered applications

The display is split into sections:

- **Summary Ribbon**
- **Discovered Application Stats**
- **Access Trend**
- **Activity charts** for *Application Accesses by Application Group*, *Most Application Accesses by Location*, *Most Application Accesses by Device*, and *Most Application Accesses by User Group*.

Each chart in the display includes a **View all** link providing access to a detail page showing log records for the corresponding chart. These log records include links to the application and user involved in the logged event. *Ivanti* recommends using this page to access the metrics page for the specific application (see Showing Usage Data for a Specific Application) or user (see Showing Activity for a Specific User). This method of navigation offers an alternative to searching for a specific application through the "Select Application" filter option, where the exact application name might not be known (for example, discovered and default applications not specifically defined in a secure access policy). To learn more about using the chart detail page, see Viewing Detailed Logs for a Chart.

The **summary ribbon** provides the following metrics:

- **All Applications**: The number of applications discovered by the *Controller*.
- **Active Applications**: The number of discovered applications accessed during the selected time period.
- **Users**: The number of users active with discovered applications.
- **Gateways**: The number of *nZTA Gateway*s serving discovered applications.
- **Anomalies**: The number of anomalous application accesses based on geographic and business hours irregularities.

The **Discovered Application Stats** panel provides two charts:

- **Top Applications Discovered by Policy**: A chart showing the application definitions, with Application Discovery enabled, for which the greatest number of applications were discovered. The segment sizes are proportional to the number of discovered applications for each application domain.
- **Top Application Protocol**: A bar chart showing the application protocol types, with Application Discovery enabled, that attracted the greatest number of application accesses during the selected time period (for example, Web, RDP, or SSH).

For both charts, hover your pointer over each bar to display a tooltip of the type and number of accesses recorded.

The **Access Trend** panel shows application access trends that occurred with discovered applications during the selected time period. You can choose to display this information through line and bar charts, or in a Sankey chart. Use the toggle icon at the top-right to select the required view:

Toggle between line/bar chart view and Sankey chart view

To expand the current view, click the Full Screen icon:

Click this icon again to return to the standard view.

In line/bar chart view. The display is split into two segments:

- A line chart showing the number of accesses for the top-10 discovered applications during each hourly period of the day
- A bar chart showing one of four data types, selected using the **Filter Bar Chart By** drop-down control:
  - Unique User Count: Shows a count of unique user activity identified during each hourly period.
  - Unique Device Type Count: Shows a count of unique device types identified during each hourly period.
  - Unique Location Count: Shows a count of activity from unique user locations identified during each hourly period.
  - Unique User Group Count: Shows a count of activity from unique user groups identified during each hourly period.

If you set a Time Period filter than spans more than one day, the data values shown in each hour period are cumulative totals for the same hour in each day during the time period.

In this chart, hover your pointer over each hourly interval to view a tooltip showing the corresponding data totals. Furthermore, you can click and drag a select box across a shorter time period to zoom in on a narrower time window. To return to the full 24 hour period, click the zoom out icon:

Zoom out from a selected time period

To toggle on or off the data for a particular application, click the name in the legend. Or, to view only the data for a specific application, click the corresponding line in the graph.

In the Sankey chart view, *nZTA* provides an alternate visualization of application access activity, showing directed flow between related objects.

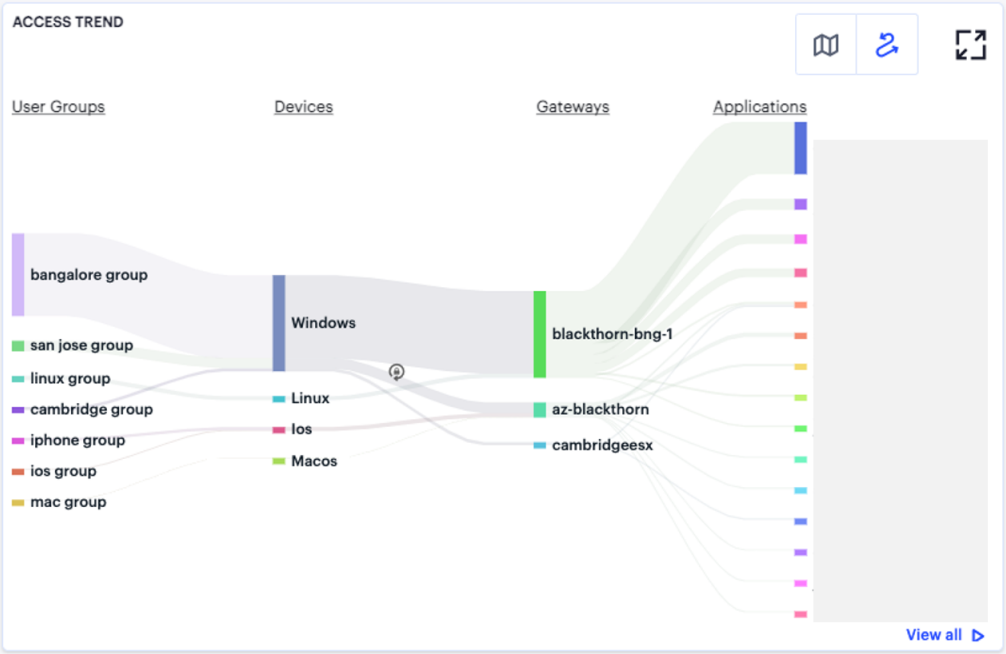

User Access Trends Sankey chart for discovered applications

The chart maps **User Groups > Devices > Gateways > Applications**. By hovering your pointer over a flow of interest, *nZTA* displays a tooltip confirming the scale of the activity between the two objects connected by the flow. To focus the display on a specific flow, or to identify related objects that interact with this flow, click the chart at a point of interest. *nZTA* provides highlighting to all flows that pass through the point selected.

The **Activity Charts** on this page represent top application access totals in the following categories:

- **Application Accesses by Application Group**: a grid chart showing the application groups containing the applications, with Application Discovery enabled, that accrued the highest number of successful accesses. Application group sizes in the chart are proportional to the number of accesses, compared with other groups. Tooltips show a count of the accesses made to that group. To learn more about Application Groups, see Adding Application Groups to the Controller.

- **Most Application Accesses by Location**: a bar chart showing a list of the most active user locations with respect to application access. Tooltips show a count of the application accesses by that location.
- **Most Application Accesses by Device**: a bar chart showing a list of the most active user device types with respect to application access. Tooltips show a count of the application accesses by that device type.
- **Most Application Accesses by User Group**: a bar chart showing a list of the most active user groups with respect to application access. Tooltips show a count of the application accesses by users in that user group.

Hover your pointer over a particular element to view a tooltip showing the label and total.

Default *ZTA Gateway* Applications

The *Controller* includes a default secure access policy called "Application discovery", disabled by default, that is used to define behavior for applications and resources that are not controlled by a specifically-created secure access policy. In this policy, you can add a default *nZTA Gateway* that you want to use to handle all such requests. To learn more about configuring a default *nZTA Gateway*, see Configuring a Default Gateway for Application Discovery.

Due to the nature of the typical background resource and API requests made by a client device during normal use of a web-based application, the metrics shown on this page might include a large number of secondary application and API requests that *nZTA* identifies and logs as not falling under the remit of the primary application's secure access policy. Such requests have been handled instead by the default *nZTA Gateway*.

> The applications listed on this tab could be operating system triggered resource requests related, for example, to the act of connecting to the internet. It should not be assumed that the URLs and IP addresses shown here are automatically connected to accessing a *nZTA*-controlled application or resource.

The *Default Gateway Applications* tab shows usage metrics for all applications and resources handled by the default *nZTA Gateway*:

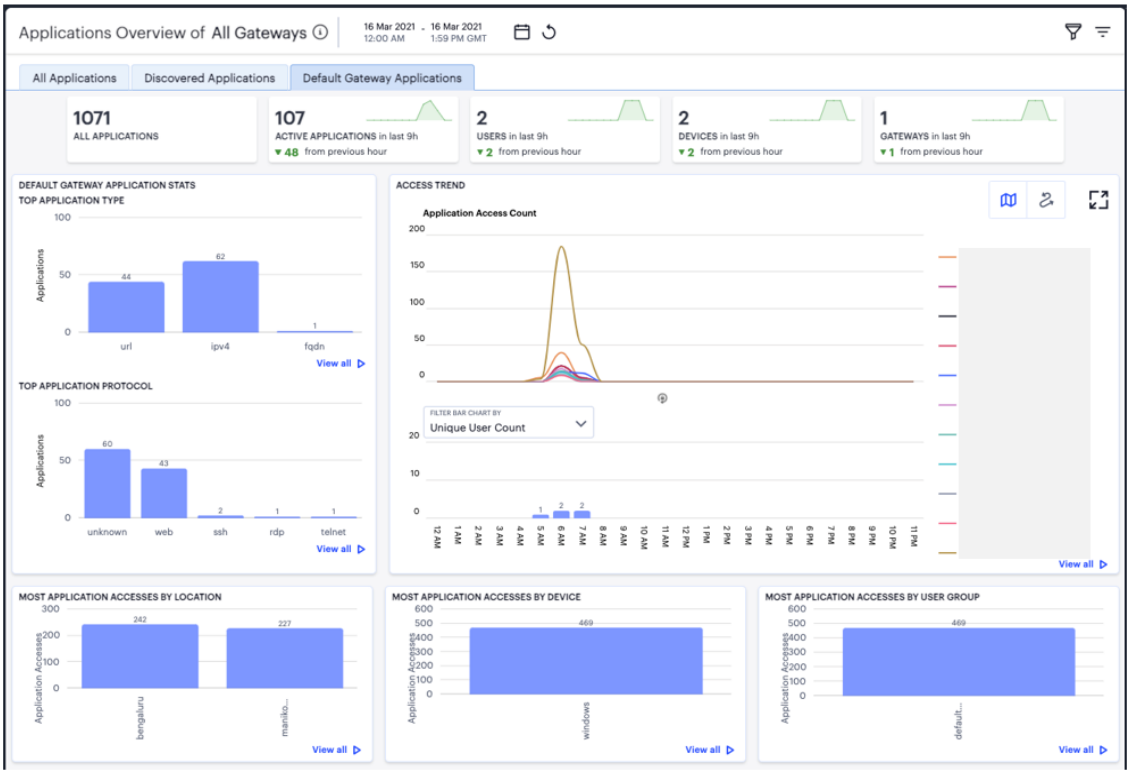

Viewing usage charts and graphs for default *nZTA Gateway* applications

The display is split into sections:

- **Summary Ribbon**
- **Default Gateway Application Stats**
- **Access Trend**
- **Activity charts** for *Most Application Accesses by Location*, *Most Application Accesses by Device*, and *Most Application Accesses by User Group*.

Each chart in the display includes a **View all** link providing access to a detail page showing log records for the corresponding chart. These log records include links to the application and user involved in the logged event. *Ivanti* recommends using this page to access the metrics page for the specific application (see Showing Usage Data for a Specific Application) or user (see Showing Activity for a Specific User). This method of navigation offers an alternative to searching for a specific application through the "Select Application" filter option, where the exact application name might not be known (for example, discovered and default applications not specifically defined in a secure access policy). To learn more about using the chart detail page, see Viewing Detailed Logs for a Chart.

The **summary ribbon** provides the following metrics:

- **All Applications**: The number of applications handled by the default *nZTA Gateway*.
- **Active Applications**: The number of default *nZTA Gateway* applications accessed during the selected time period.
- **Users**: The number of users active with default *nZTA Gateway* applications.
- **Devices**: The number of devices accessing default *nZTA Gateway* applications.
- **Gateways**: Denotes the *nZTA Gateway*, or number of *nZTA Gateways* in the Gateway Group, selected as the *default Gateway* in the "Application Discovery" secure access policy.

The application details shown here are unique to this page and are not included in other summary ribbons or metrics involving *all* applications.

The **Default Gateway Application Stats** panel provides two charts:

- **Top Application Type**: A bar chart showing the application types that attracted the greatest numbers of application accesses during the selected time period (for example, FQDN, URL, or IP address).
- **Top Application Protocol**: A bar chart showing the application protocol types that attracted the greatest number of application accesses during the selected time period (for example, Web, RDP, or SSH).

For both charts, hover your pointer over each bar to display a tooltip of the type and number of accesses recorded.

The **Access Trend** panel shows application access trends that occurred during the selected time period. You can choose to display this information through line and bar charts, or in a Sankey chart. Use the toggle icon at the top-right to select the required view:

Toggle between line/bar chart view and Sankey chart view

To expand the current view, click the Full Screen icon:

Expand the current view

> Click this icon again to return to the standard view.

In line/bar chart view. The display is split into two segments:

- A line chart showing the number of accesses for the top-10 requested applications during each hourly period of the day
- A bar chart showing one of four data types, selected using the **Filter Bar Chart By** drop-down control:
  - Unique User Count: Shows a count of unique user activity identified during each hourly period.
  - Unique Device Type Count: Shows a count of unique device types identified during each hourly period.
  - Unique Location Count: Shows a count of activity from unique user locations identified during each hourly period.
  - Unique User Group Count: Shows a count of activity from unique user groups identified during each hourly period.

> If you set a Time Period filter than spans more than one day, the data values shown in each hour period are cumulative totals for the same hour in each day during the time period.

In this chart, hover your pointer over each hourly interval to view a tooltip showing the corresponding data totals. Furthermore, you can click and drag a select box across a shorter time period to zoom in on a narrower time window. To return to the full 24 hour period, click the zoom out icon:

Zoom out from a selected time period

To toggle on or off the data for a particular application, click the name in the legend. Or, to view only the data for a specific application, click the corresponding line in the graph.

In the Sankey chart view, *nZTA* provides an alternate visualization of application access activity, showing directed flow between related objects.

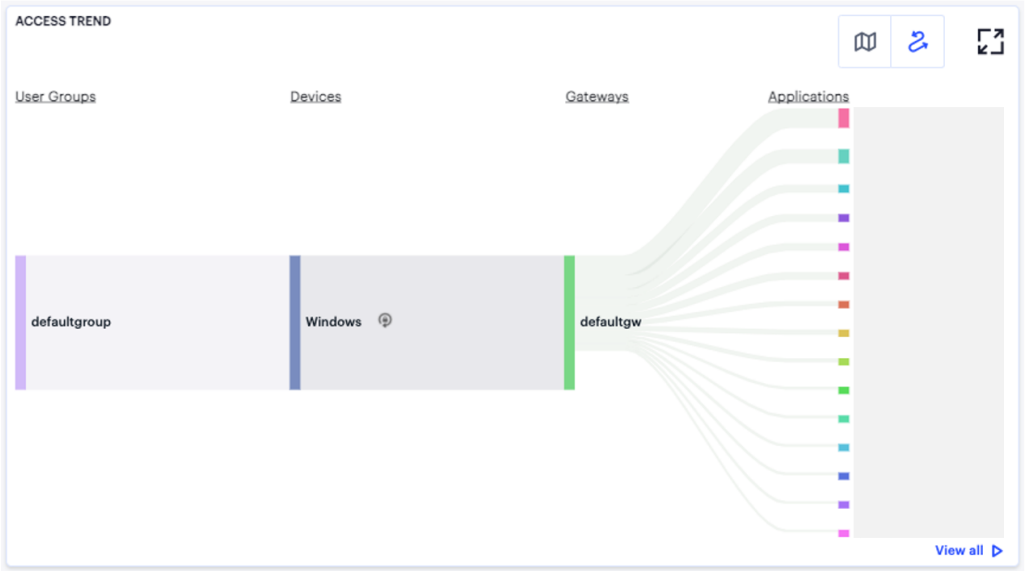

User Access Trends Sankey chart for default *nZTA Gateway* applications

The chart maps **User Groups > Devices > Gateways > Applications**. By hovering your pointer over a flow of interest, *nZTA* displays a tooltip confirming the scale of the activity between the two objects connected by the flow. To focus the display on a specific flow, or to identify related objects that interact with this flow, click the chart at a point of interest. *nZTA* provides highlighting to all flows that pass through the point selected.

The **Activity Charts** on this page represent top application access totals in the following categories:

- **Most Application Accesses by Location**: a bar chart showing a list of the most active user locations with respect to application access. Tooltips show a count of the application accesses by that location.
- **Most Application Accesses by Device**: a bar chart showing a list of the most active user device types with respect to application access. Tooltips show a count of the application accesses by that device type.
- **Most Application Accesses by User Group**: a bar chart showing a list of the most active user groups with respect to application access. Tooltips show a count of the application accesses by users in that user group.

Hover your pointer over a particular element to view a tooltip showing the label and total.

When viewing metrics on this page, the following limitations should be noted:

- Non-Compliance messages are not generated for the default *nZTA Gateway*. This is due to the fact *Ivanti Secure Access Client* blocks such messages directly without sending them on to the *nZTA Gateway*.
- The default *nZTA Gateway* application details captured on this page are not included in the metrics captured on the *Network Overview* page. However, the default *nZTA Gateway* is still shown on the *Network Overview* page for monitoring purposes (for example, CPU, disk, and memory usage).
- Anomaly detection for applications handled by the default *nZTA Gateway* (especially business hours anomalies) is not displayed on any of the *Insights* dashboards. This is due to the fact that the number of applications detected can be very large, which can in turn impact the user UEBA Threat score.
- Log records for applications handled by the default *nZTA Gateway* are displayed only on the **Secure Access > Gateways > Logs** page. These records are not displayed on the **Insights > Logs** page.
- If a user associated with the default User Group tries to access applications handled by the default *nZTA Gateway*, the *Users* dashboard for that specific user displays only the user UEBA Threat score, risk rank and the moving average on the *Access Trend* chart. It does not capture details of the default gateway applications accessed, primarily to ensure that the application data displayed here does not become overpopulated.

Configuring Default Gateway Application

A Configure button is provided in the *nZTA Gateway* Applications page to trigger the workflow of blocking the discovered applications behind default gateway.

To configure default *nZTA Gateway* application:

1. In the default *nZTA Gateway* applications page, click **Configure**.

   The *Configure Applications* page is displayed showing a list of discovered applications behind the default gateway.

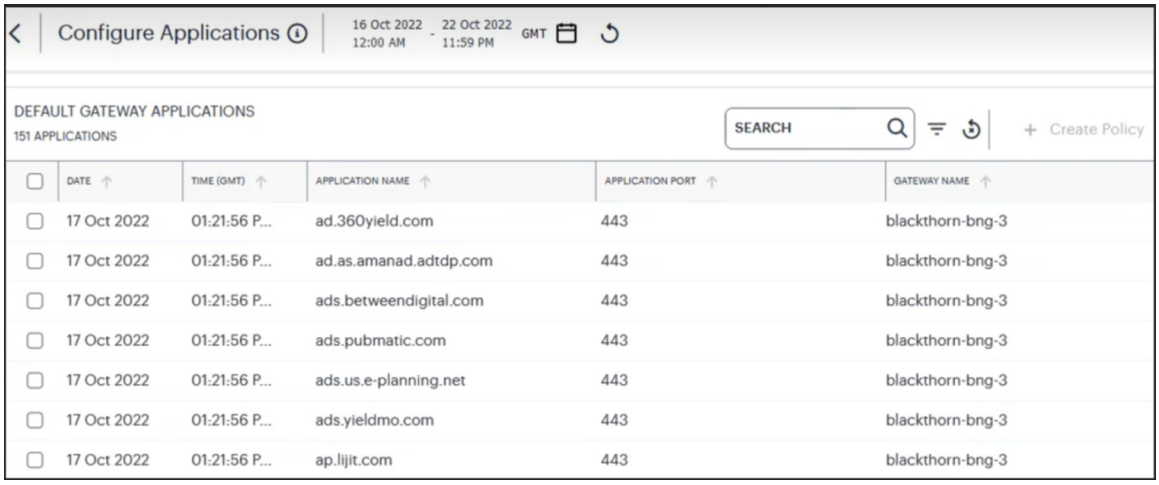

   Configure default *nZTA Gateway* applications

2. In the search box provided, start typing the application name. nZTA auto-completes any matching application name.

3. Select an application from the list and click **Create Policy** to create a Secure Access Policy. To learn more about creating a secure access application, see Creating/Editing Secure Access Policies.

Showing Usage Data for a Specific Application

To view usage data and metrics for a specific application, use the **Set Application** option in the filter menu:

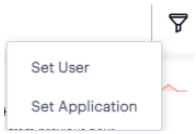

Activating the Set Application option

Alternatively, from the **Network Overview** page, access specific application data from the **Applications** info-panel view. For more details, see Using the Summary Ribbon.

nZTA displays the *Application* page, showing activity for the selected application:

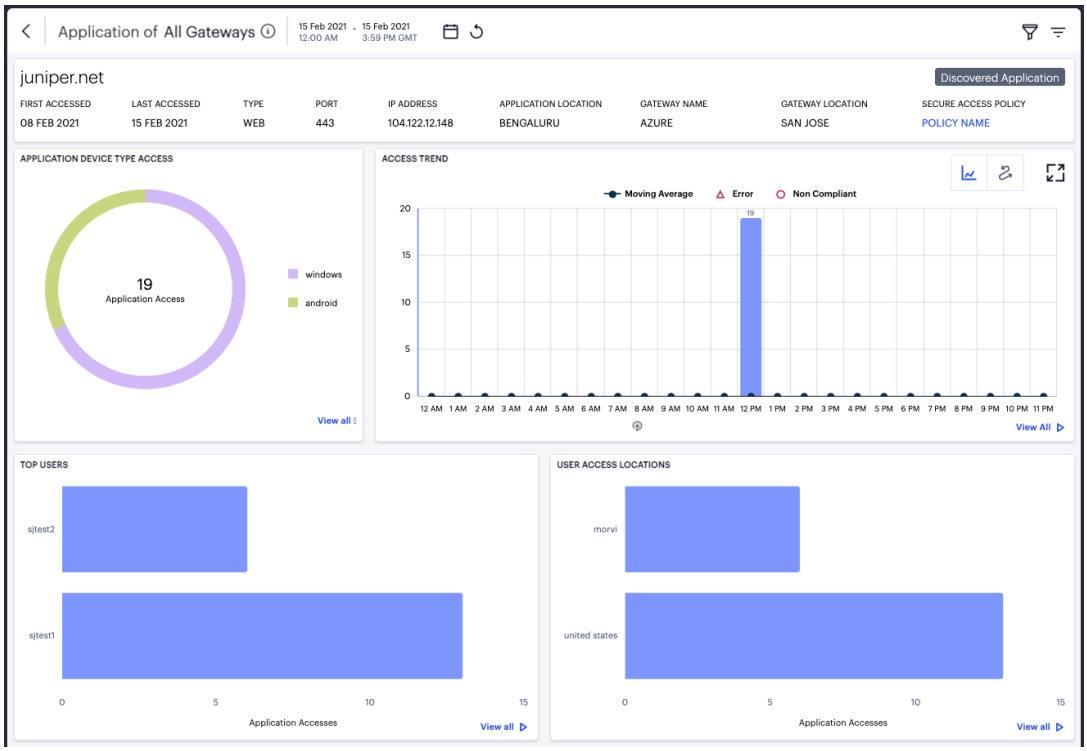

Viewing usage data for a specific application

Understanding the Display

The *Application* page contains the following components:

- **Filter bar**, allowing the selection of active or historic data. For details, see Using the Filter Bar.
- **Summary ribbon**, showing details of the selected application. For more details, see Using the Summary Ribbon.
- **Application Device Type Access**, showing accesses per device type. For more details, see Viewing Application Accesses by Device Type.
- **Access trend**, showing application accesses, non-compliance, and errors over time. For more details, seeViewing Access Trends for the Selected Application.
- **Activity charts**, showing top users and locations. For more details, see Viewing Application Activity Charts.

Each chart on this page includes a **View all** link. This link provides access to a detail view showing logs for the corresponding chart. For example:

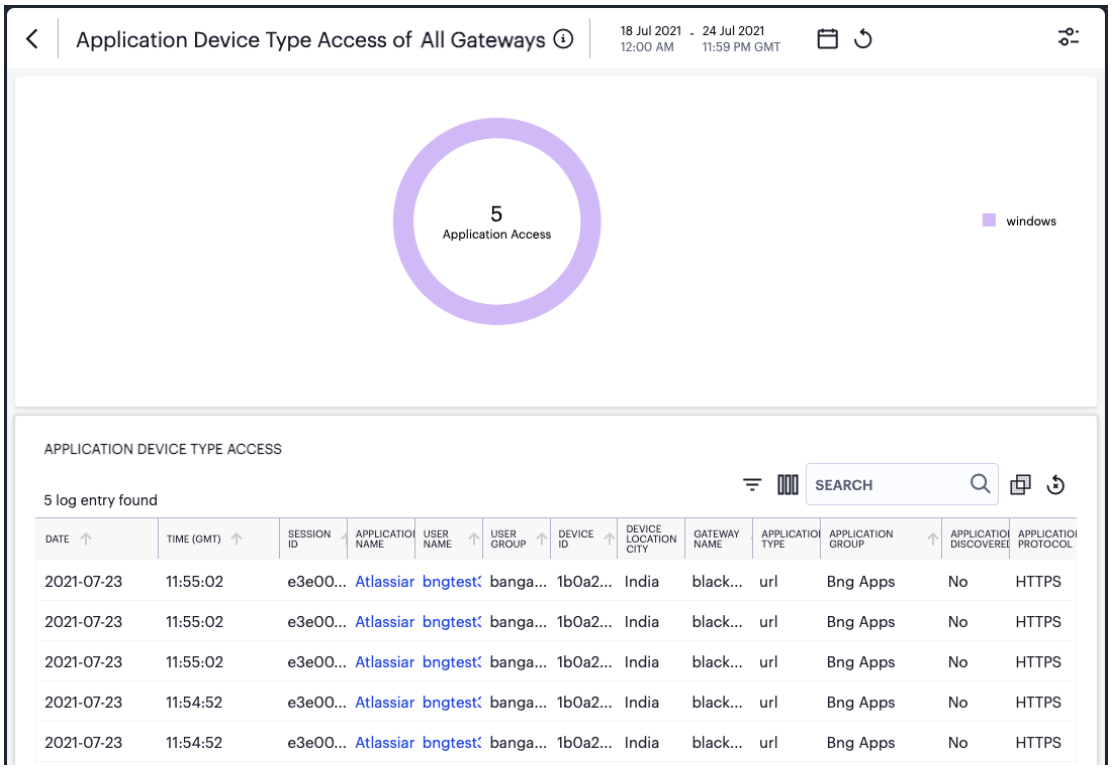

Viewing Application Device Type Access detailed logs

Each detail view shows logs for the corresponding chart or category. To learn more about using the chart detail page, see Viewing Detailed Logs for a Chart.

Using the Summary Ribbon

The Summary Ribbon at the top of the *Application* page shows details for the application:

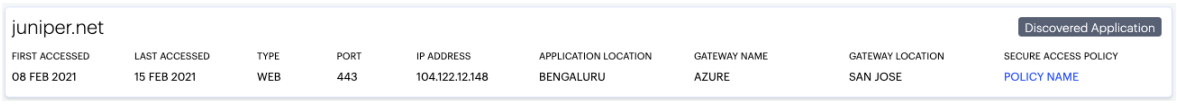

Viewing the summary ribbon

The summary ribbon provides the following information:

- **First Accessed**: The date on which the application was first accessed.
- **Last Accessed**: The date on which the application was most recently accessed.

- **Type**: The application type. For example, "Web".
- **Port**: The port on which the application is accessed.
- **IP Address**: The IP address through which the application is accessed.
- **Application Location**: The geographic location where the application is hosted.
- **Gateway Name**: The name of the *nZTA Gateway* managing the application.
- **Gateway Location**: The location of the *nZTA Gateway* managing the application.
- **Secure Access Policy**: The name of the Secure Access Policy governing access to the application. Click the name of the policy to access the *Secure Access Policies* page.

> If your application is **discovered**, this is denoted by a label in the ribbon. To learn more about discovered applications, see Reviewing Application Usage.

Viewing Application Accesses by Device Type

The *Application Device Type Access* panel shows application accesses by device type:

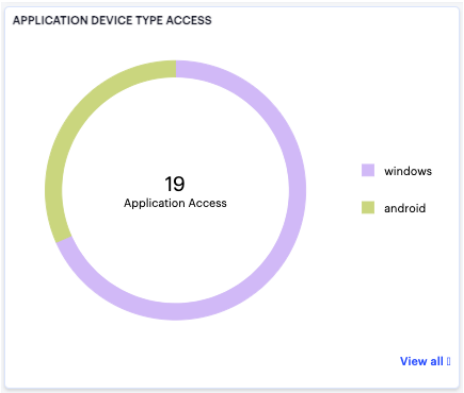

Viewing application accesses by device type

The chart provides a breakdown of application accesses for each device type. The number in the center of the chart is a total for all device types. Hover your pointer over a device type to view a tooltip showing the number of accesses made by devices of that type.

Viewing Access Trends for the Selected Application

nZTA uses this section to show access trends for the selected application that occurred during the selected time period.

You can choose to display this information through line and bar charts (as shown), or in a Sankey chart. Use the toggle icon at the top-right to select the required view:

Toggle between line/bar chart view and Sankey chart view

To expand the current view, click the Full Screen icon:

Expand the current view

The line and bar chart shows application access trends through each hourly interval of the day:

Viewing hourly access trends for an application

The horizontal axis reflects the 24 hourly intervals in a day, with the vertical axis showing the number of accesses. As users access the application, access attempts are recorded and shown in this chart according to the hour in which they occurred. This in turn provides an overview of the daily access trends for the application.

The bars denote the number of accesses made, with the line charting the moving average over 30 days.

Errors and non-compliances detected in the hour are marked as per the legend.

Hover your pointer over an hour interval to see a tooltip summary of data points for that hour:

Viewing a data summary tooltip

To learn more about the application access Sankey chart, see Viewing the Application Access Trends Sankey Chart.

Viewing the Application Access Trends Sankey Chart

The Application Access Trends Sankey chart provides an alternate visualization of access activity for your selected application, showing directed flow between related objects. The width of each stream in the flow is proportional to the utilization of the object the flow passes through, allowing an administrator to view significant usage trends for the selected application. `

To toggle between the application access trends line/bar chart and the application access trends Sankey chart, use the icons at the top-right:

Toggle between line/bar chart view and Sankey chart view

By clicking the toggle display icon, the Application Access Trends Sankey chart replaces the line/bar chart in the display. All other components remain unchanged.

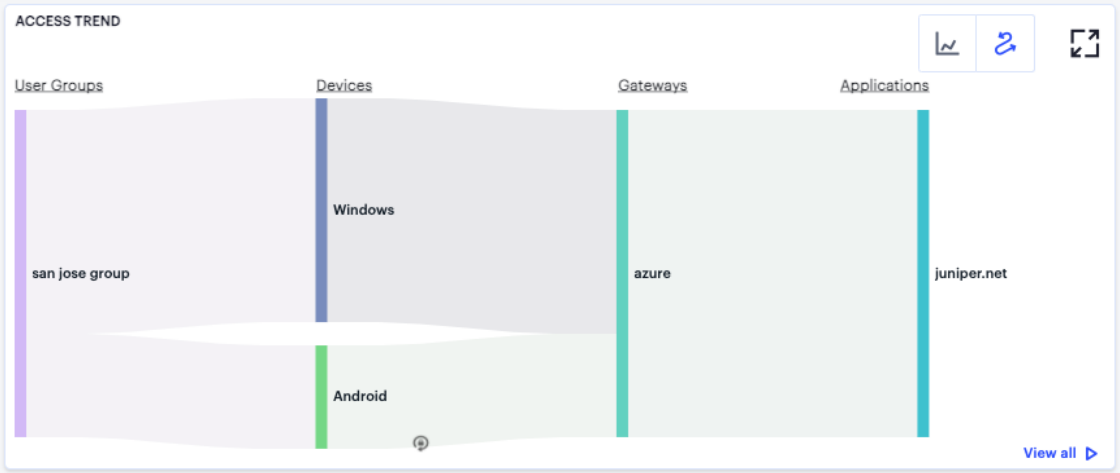

Displaying the Application Access Sankey Chart

The Sankey chart maps **Locations** > **Devices** > **Gateways** > **Application** for the selected application. By hovering your pointer over a flow of interest, *nZTA* displays a tooltip confirming the scale of the activity between the two objects connected by the flow.

To focus the display on a specific flow, or to identify related objects that interact with this flow, click the chart at a point of interest. *nZTA* provides highlighting to all flows that pass through the point selected.

Viewing Application Activity Charts

On the *Application* page, *nZTA* provides the following charts:

- **Top Users**: Shows the users who accrued the most accesses for the selected application.
- **User Access Locations**: Shows the user locations from which the most accesses were recorded for the selected application.

Hover your pointer over a bar in the chart to view a tooltip showing the number of accesses made.

Viewing Currently Enrolled User Devices

You can view a list of the user devices currently enrolled with the *Controller* through the Tenant Admin Portal.

To view currently enrolled user devices:

1. Log into the Tenant Admin Portal.
2. Click **Insights** > **Devices**.

The **Devices** page appears.

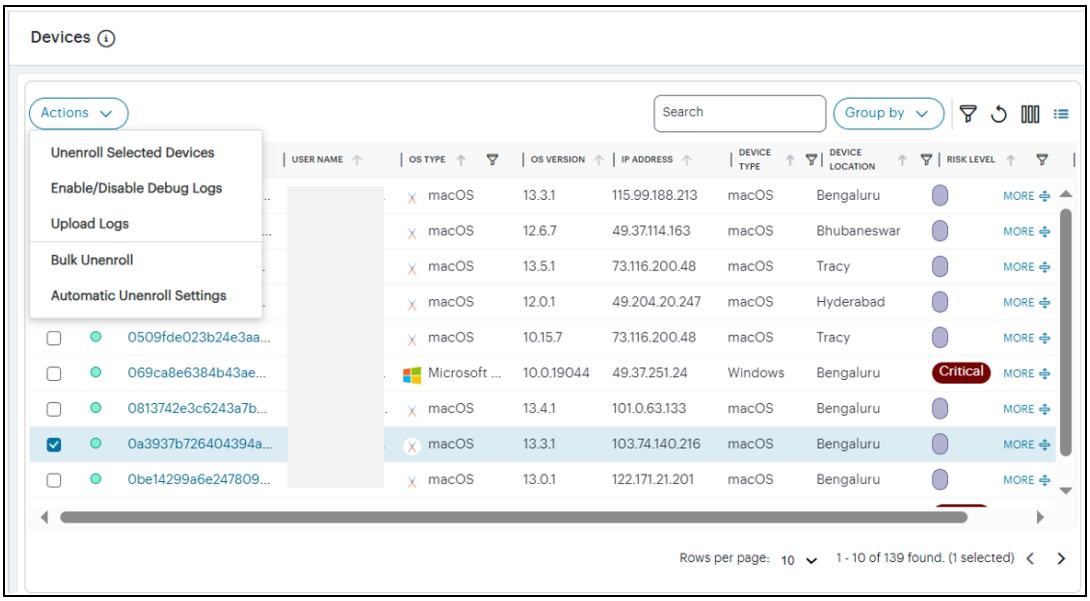

Enrolled user device list

Use this page to:

- View information concerning the devices your users have enrolled to your *nZTA* service.
- Group the records by unique data item using the Group By drop-down list.
- Filter the records based on column data using the column filters.
- Perform following using the **Actions** drop-down menu:
    - Unenroll selected devices.
    - Enable/Disable debug logs.
    - Upload *Ivanti Secure Access Client* log files from selected devices to the *Controller* for analysis, see Uploading Client Logs to the Controller.
    - Do a Bulk enroll.
    - Set Automatic unenroll.

If you click Device ID link within a specific row from the table of devices, *nZTA* displays an info-panel providing further details:

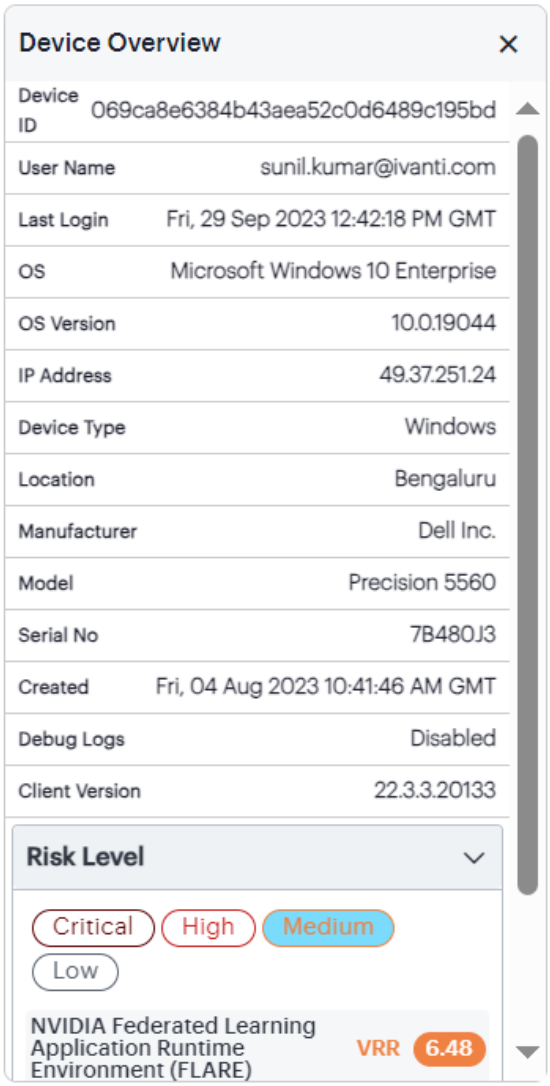

Enrolled user device info-panel

To unenroll selected devices from the list, tick the check box in the row for one or more chosen devices, click the **Actions** drop-down list, and then select **Unenroll Selected Devices**.

To automatically unenroll inactive devices after a set time period, click the **Actions** drop-down list, and then select **Automatic Unenroll Settings**.

Then, in the *Automatically Unenroll Devices* dialog, select an inactivity duration and click **Save**. All devices that reach this duration limit without being logged-in are automatically unenrolled. To disable the automatic unenroll mechanism, disable **Automatic Unenroll Enabled** switch.

Uploading Client Logs to the *Controller*

To help *Ivanti* Technical Support teams troubleshoot and debug any issues your end user devices have when connecting to your *nZTA* services, you can upload logs from connected devices to the *Controller* for analysis. This process can be initiated remotely from the *nZTA* Tenant Admin portal, or from the *Ivanti Secure Access Client* application installed on the device.

Client connection logs are stored securely in the *Controller* and cannot be accessed by unauthorized users. To obtain the logs, your device must be in a *connected* state.

> Upload of client logs is not currently supported on *Ivanti Secure Access Client* Linux variants.

To upload *Ivanti Secure Access Client* logs for an enrolled device from the *nZTA* Tenant Admin Portal:

1. Log into the Tenant Admin Portal.
2. Click **Secure Access > Devices > Enrolled Devices**.

   The **Enrolled Devices** page appears.
3. From the list of enrolled devices, tick the check box in the row for your selected device and click the following icon:

   

   Upload client logs for the selected device

> Devices running *Ivanti Secure Access Client* versions released earlier than *nZTA* 20.11 do not have the upload capability and cannot respond to this feature.

> If your device is enrolled, but not connected, an error message is displayed "User session is not found".

4. *nZTA* displays a confirmation dialog showing the details of the task:

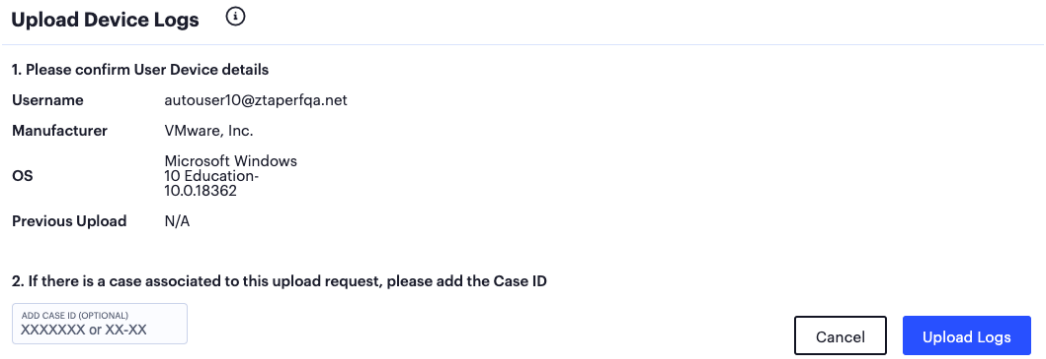

Confirming upload of the selected client logs

5. Confirm the details in *section 1* and provide the optional **case ID** (if instructed to do so by your support representative) in *section 2*.

6. To initiate the upload process, click **Upload Logs**.

   The upload process begins.

7. A log upload status section is added to the info-panel for that device. This shows the progress of the upload and includes a timestamp of the last upload:

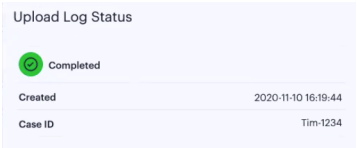

Uploaded log status for a device - in progress

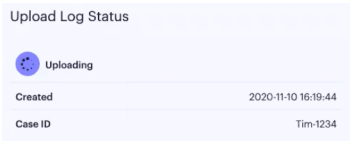

Uploaded log status for a device - complete

To upload logs from the end user device:

1. On the end user device, open the *Ivanti Secure Access Client* application.

2. In the *Ivanti Secure Access Client* application, use the **Upload** facility in the **Logs** sub-menu.

For example:

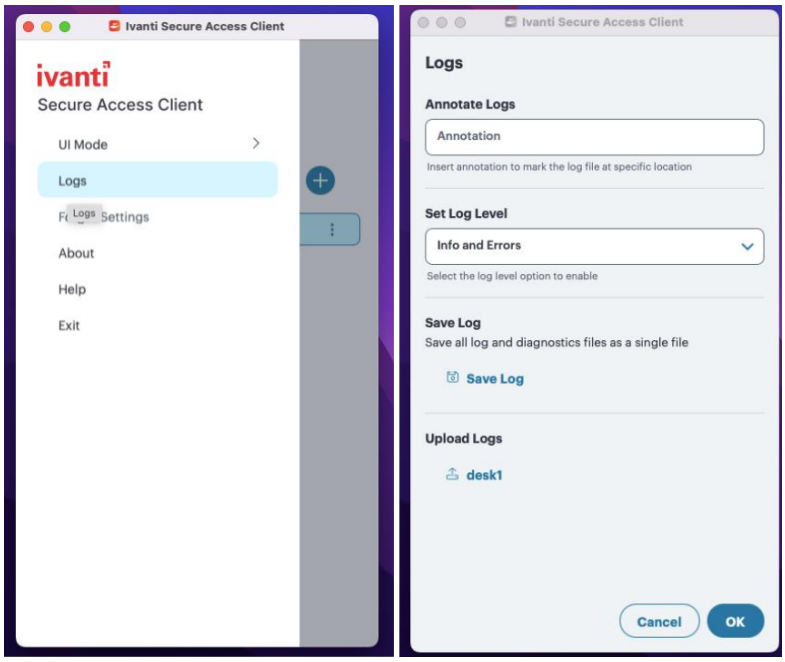

Uploading logs

> The log upload facility within *Ivanti Secure Access Client* requires the device to be enrolled with the *Controller*. However, a user can initiate a log upload in both connected and disconnected states.

Monitoring *nZTA Gateway* Activity

To view usage data and metrics for all *nZTA Gateway*s, or for a specific *nZTA Gateway*, use the *Gateways Overview* page.

To view the *Gateways Overview* page, select **Insights > Gateways**:

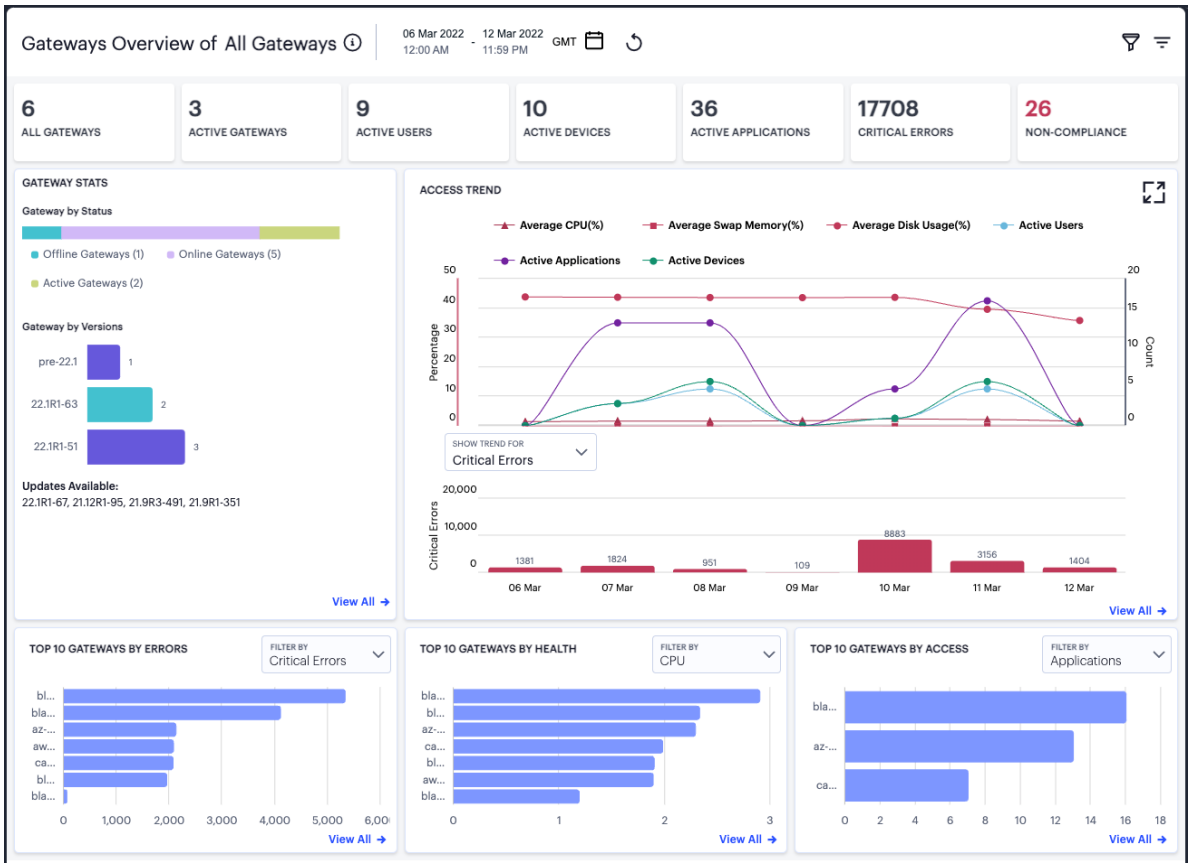

Viewing *nZTA Gateway* metrics

By default, this page shows data for all *nZTA Gateways*. To view data for a specific *nZTA Gateway*, use the filter feature described in Using the Filter Bar.

> Some features on this page require your *nZTA Gateways* to be running as version 22.1R1 or later. *nZTA Gateways* running versions earlier than this might not be included in some status and health data.

Understanding the Display

The *Gateways Overview* page contains the following components:

- **Filter bar**, allowing the selection of active or historic data. For details, see Using the Filter Bar.
- **Summary ribbon**, showing metrics for deployed *nZTA Gateways*. For more details, see Using the Summary Ribbon.

- **Gateway Stats**, showing an overview of the status of your deployed *nZTA Gateways*. For more details, see Reviewing the Status of your Deployed nZTA Gateways.
- **Access trend**, showing *nZTA Gateway* usage metrics over time. For more details, see Viewing nZTA Gateway Access Trends.
- **Activity charts**, showing top 10 *nZTA Gateway* usage metrics in a number of categories. For more details, see Viewing nZTA Gateway Activity Charts.

Each chart on this page includes a **View all** link. This link provides access to a detail view showing logs for the corresponding chart. For example:

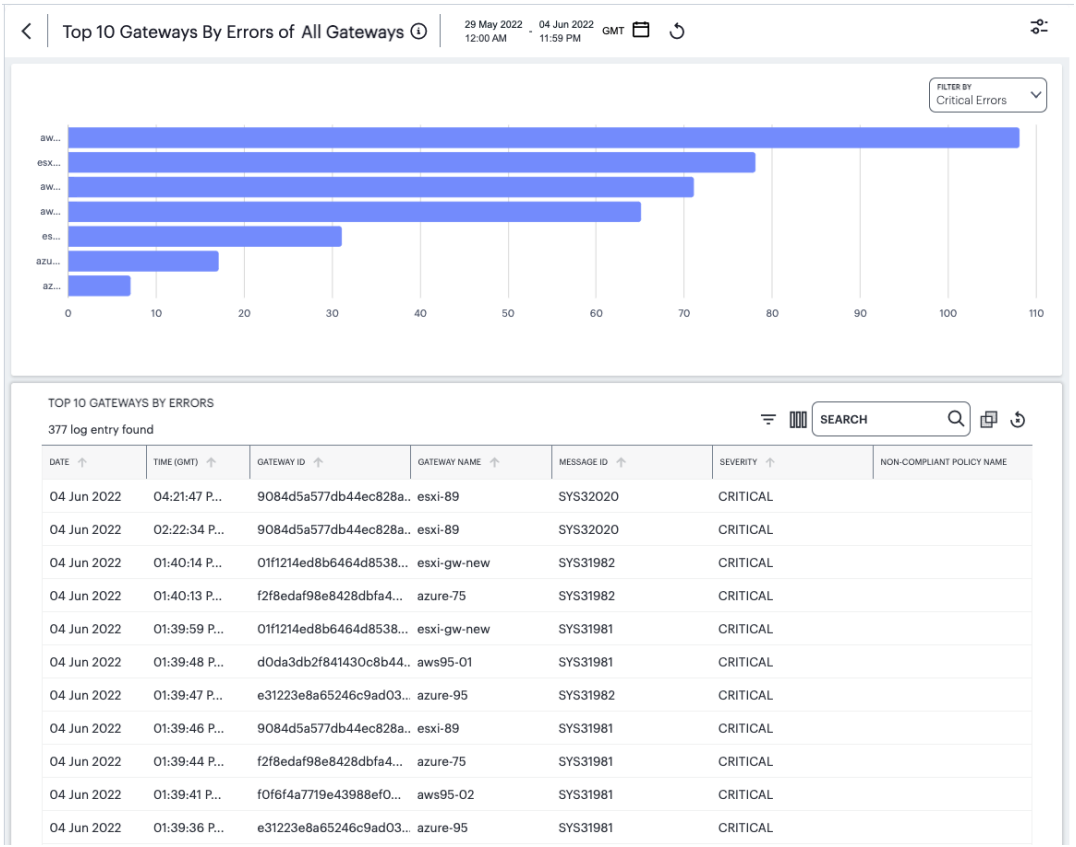

Viewing detailed logs for Top 10 Gateways by Errors

Each detail view shows logs for the corresponding chart or category. To learn more about using the chart detail page, see Viewing Detailed Logs for a Chart.

Using the Summary Ribbon

The Summary Ribbon at the top of the *Gateways Overview* page shows relevant summary statistics relating to your deployed *nZTA Gateway*s:

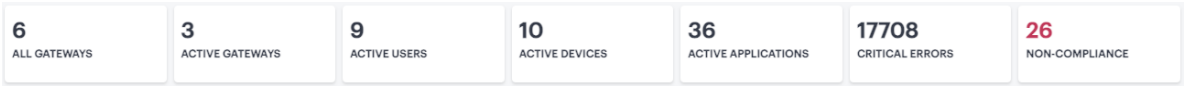

Viewing the summary ribbon

The summary ribbon provides the following information:

- **All Gateways**: The total number of deployed *nZTA Gateway*s.
- **Active Gateways**: The number of active deployed *nZTA Gateway*s in the selected time period. In other words, those *nZTA Gateway*s that are online and reporting activity.
- **Active Users**: The number of users accessing applications and resources managed by your active *nZTA Gateway*s during the selected time period.
- **Active Devices**: The number of unique devices used to access applications and resources managed by your active *nZTA Gateway*s during the selected time period.
- **Critical Errors**: The number of critical errors observed on your *nZTA Gateway*s during the selected time period.
- **Non-Compliance**: The number of non-compliant attempts to access the applications managed by your *nZTA Gateway*s.

Reviewing the Status of your Deployed *nZTA Gateway*s

The **Gateway Stats** panel shows the status of your deployed *nZTA Gateway*s during the selected time period:

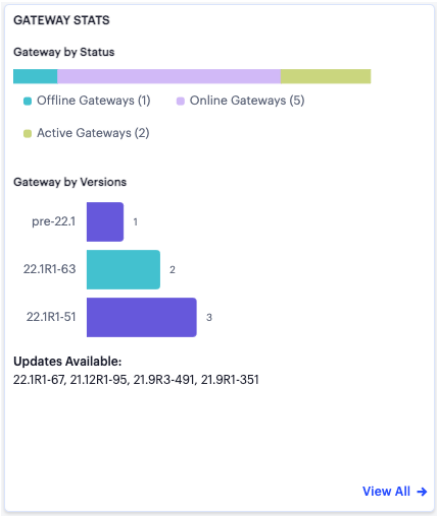

Viewing the Status of your deployed *nZTA Gateway*s

The panel provides three separate components:

- A visual breakdown of your deployed *nZTA Gateway*s as *Offline Gateways* (*nZTA Gateway*s deployed but unresponsive/unavailable), *Online Gateways* (responsive *nZTA Gateway*s not currently handling user traffic), and *Active Gateways* (*nZTA Gateway*s handling user traffic).
- A visual breakdown of your deployed *nZTA Gateway*s by version.
- The software updates available for your deployed *nZTA Gateway*s, if applicable.

Viewing *nZTA Gateway* Access Trends

This section shows *nZTA Gateway* access trends that occurred during the selected time period:

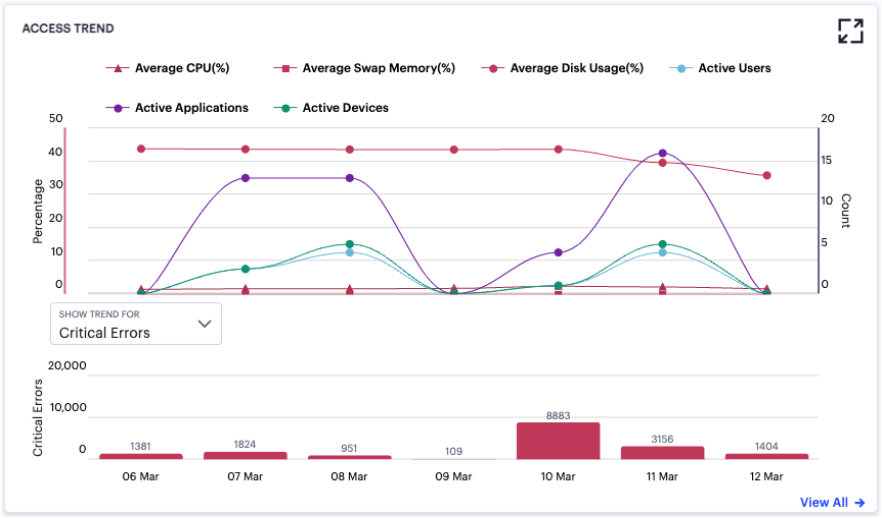

Viewing *nZTA Gateway* access trends

The horizontal axis of each chart reflects the selected time period, and dynamically adapts to span the period in increments appropriate to that period. For example, 5 minute intervals for the *Last Hour* view, or hourly intervals where you select a whole day.

To expand the current view, click the Full Screen icon:

Expand the current view

The display is split into two segments:

- A line chart showing the number of *nZTA Gateway* accesses in the selected time period.
- A bar chart showing access trends for a selected data type. Use the **Show Trend For** drop-down control to select the chart data type. Choose from:
    - Critical Errors
    - Throughput (MB)
    - Major Errors
    - Non Compliances

In this chart, hover your pointer over each interval point to view a tooltip showing the corresponding data totals. Furthermore, you can click and drag a select box across a shorter time period to zoom in on a narrower time window. To return to the full period, click the *zoom out* icon:

Zoom out from a selected time period

In the line chart, toggle on or off the data for a particular trend type by clicking the name of the type in the legend.

Viewing *nZTA Gateway* Activity Charts

On the *Gateways Overview* page, *nZTA* provides the following charts:

- **Top 10 Gateways by Errors**: The top 10 *nZTA Gateway*s for which errors were reported. Use the **Filter By** drop-down control to select the criteria for the chart. Choose from *Critical Errors*, *Major Errors*, or *Non-Compliances*.
- **Top 10 Gateways by Health**: The top 10 *nZTA Gateway*s by system health. Use the **Filter By** drop-down control to select the criteria for the chart. Choose from average *CPU* usage, average *Swap Memory* usage, average *Disk Usage*, or *Network Throughput*.
- **Top 10 Gateways by Access**: The top 10 *nZTA Gateway*s by the number of accesses. Use the **Filter By** drop-down control to select the criteria for the chart. Choose from *Applications*, *Users*, or *Devices*.

Hover your pointer over a bar in the chart to view a tooltip showing the *nZTA Gateway* name and total applicable to that bar.

Reviewing Policy Failures

When a device attempts to access an application or resource controlled by a Secure Access Policy, the device must first comply with all relevant *device policies*. If the device does not meet one or more of the conditions in a policy, a failure event is recorded and access is denied. *nZTA* displays policy failure data and metrics in the *Policy Failures* page.

To view the *Policy Failures* page, select **Insights > Policy Failures**:

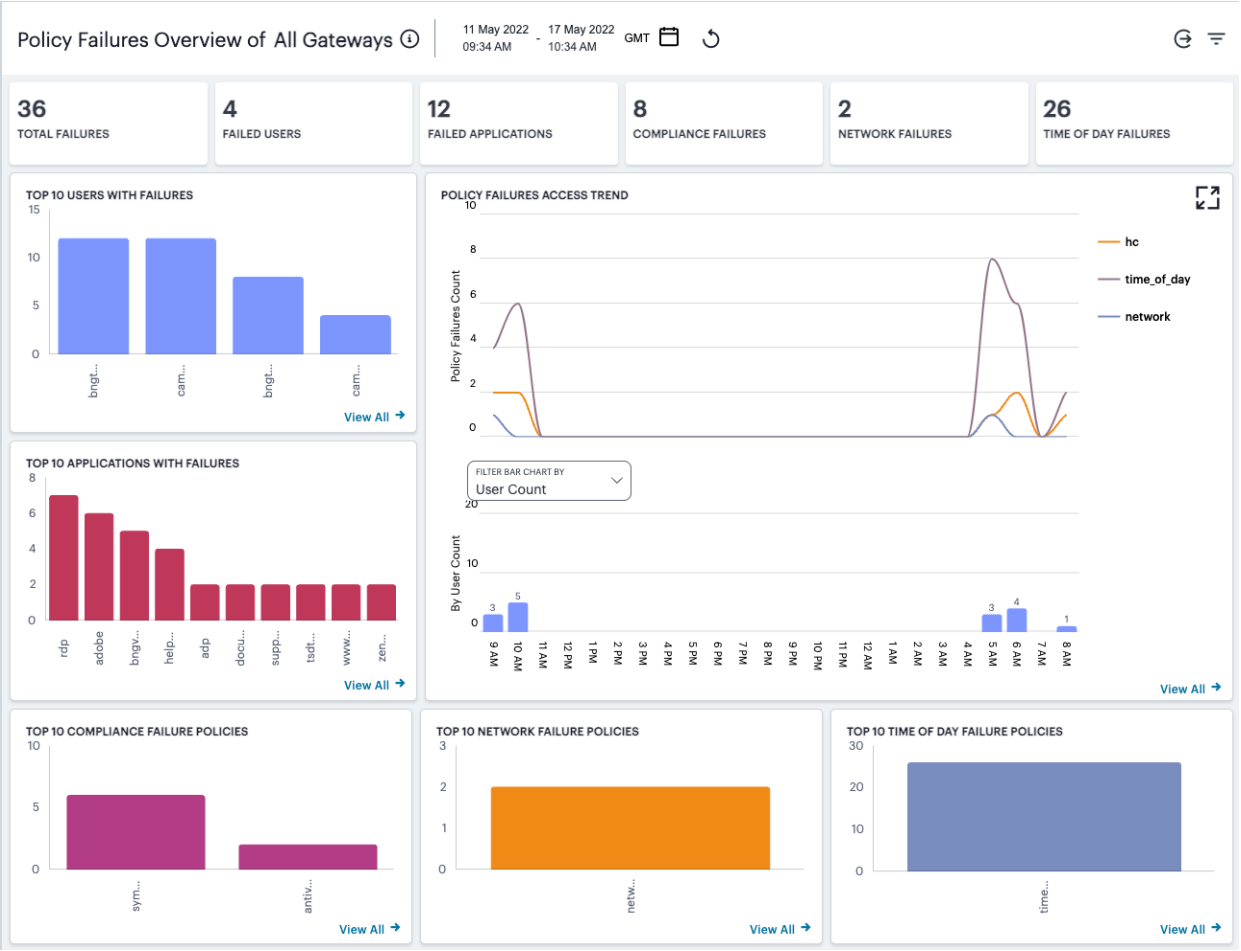

Viewing policy failure metrics

The failure types reported on this page are comprised of the following types:

- Network policy failures: a device does not meet the conditions in a policy containing a *Network* type device rule.
- Time-of-day policy failures: a device does not meet the conditions in a policy containing a *Time of day* type device rule.
- Compliance policy failures: a device does not meet the conditions in a policy containing other device compliance rules.

> In this release, policy failures based on rules of type *Location* are not included in these metrics.

For more information on configuring device rules and policies, see Creating Device Policies and Device Policy Rules.

Understanding the Display

The *Policy Failures* page contains the following components:

- **Filter bar**, allowing the selection of active or historic data. For details, see Using the Filter Bar.
- **Summary ribbon**, showing metrics pertaining to detected policy failures across various categories. For more details, see Using the Summary Ribbon.
- **Policy Failures Access trend**, showing policy failure counts over time. For more details, see Viewing Policy Failures Access Trends.
- **Activity charts**, showing top 10 failure counts in various categories. For more details, see Viewing Policy Failure Activity Charts.

Each chart on this page includes a **View all** link. This link provides access to a detail view showing logs for the corresponding chart. For example:

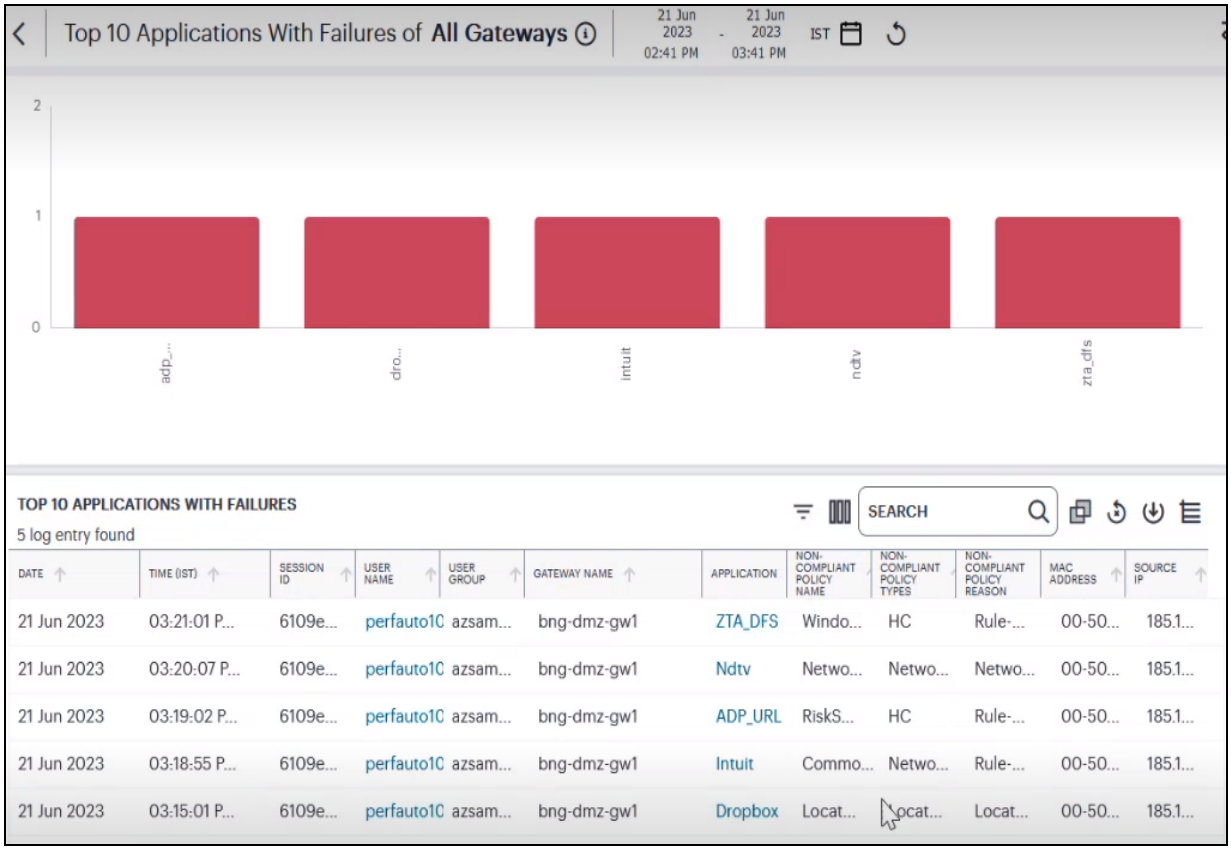

Viewing detailed logs for Top 10 Applications With Failures

Each detail view shows logs for the corresponding chart or category. To learn more about using the chart detail page, see Viewing Detailed Logs for a Chart.

Using the Summary Ribbon

The Summary Ribbon at the top of the *Policy Failures* page shows policy failure totals across a number of categories:

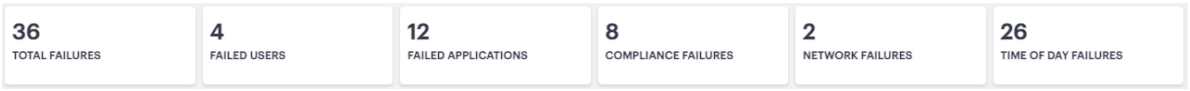

Viewing the summary ribbon

The summary ribbon provides the following information:

- **Total Failures**: The total number of policy failures detected across your deployment in the selected time period.
- **Failed Users**: The number of users who triggered a policy failure upon attempting to access an application or resource controlled by a Secure Access Policy.
- **Failed Applications**: The number of applications to which access was denied due to a policy failure.
- **Compliance Failures**: The number of compliance failures recorded against all device policies, excluding network and time-of-day type policies.
- **Network Failures**: The number of failures recorded against a network type device policy.
- **Time of Day Failures**: The number of failures recorded against a time-of-day type device policy.

Viewing Policy Failures Access Trends

nZTA uses this section to show policy failure access trends that occurred during the selected time period:

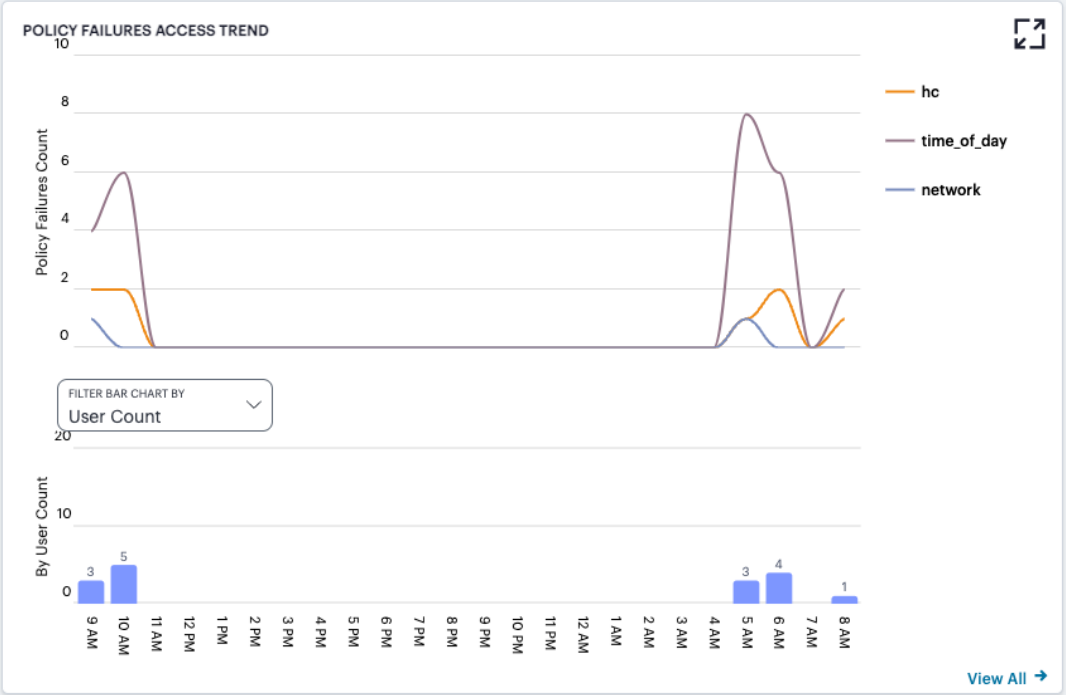

Viewing policy failure access trends

To expand the current view, click the Full Screen icon:

Expand the current view

The display is split into two segments:

- A line chart showing the number of policy failures for network, time-of-day, and compliance (marked in the chart as "hc") policy types during each hourly period of the day
- A bar chart showing one of two data types, selected using the **Filter Bar Chart By** drop-down control:
    - User Count
    - Application Count

If you set a *Time Period* filter than spans more than one day, the data values shown in each hour period are cumulative totals for the same hour in each day during the time period.

In this chart, hover your pointer over each hourly interval to view a tooltip showing the corresponding data totals. Furthermore, you can click and drag a select box across a shorter time period to zoom in on a narrower time window. To return to the full 24 hour period, click the *zoom out* icon:

Zoom out from a selected time period

In the Policy Failures Count line chart, toggle on or off the data for a particular failure trend type by clicking the name of the type in the legend. Or, to view only the data for a specific type, click the corresponding line in the graph.

Viewing Policy Failure Activity Charts

On the *Policy Failures* page, *nZTA* provides the following charts:

- **Top 10 Users With Failures**: The top 10 users users who triggered a policy failure upon attempting to access an application or resource controlled by a Secure Access Policy.
- **Top 10 Applications With Failures**: The top 10 applications to which access was denied due to a policy failure.
- **Top 10 Compliance Failure Policies**: The top 10 compliance device policies that reported failures.
- **Top 10 Network Failure Policies**: The top 10 network device policies that reported failures.
- **Top 10 Time of Day Failure Policies**: The top 10 time-of-day device policies that reported failures.

Hover your pointer over a bar in the chart to view a tooltip showing the number of failure in that case.

Checking the Logs

The *nZTA* Logs page displays audit and activity events observed by your *nZTA* secure access infrastructure. These events are reported to the *Controller* by your *nZTA Gateways* and the Authentication, Authorization and Accounting (AAA) service.

To view the Logs page:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Insights** icon, then select **Logs**.

The *Logs* page appears.

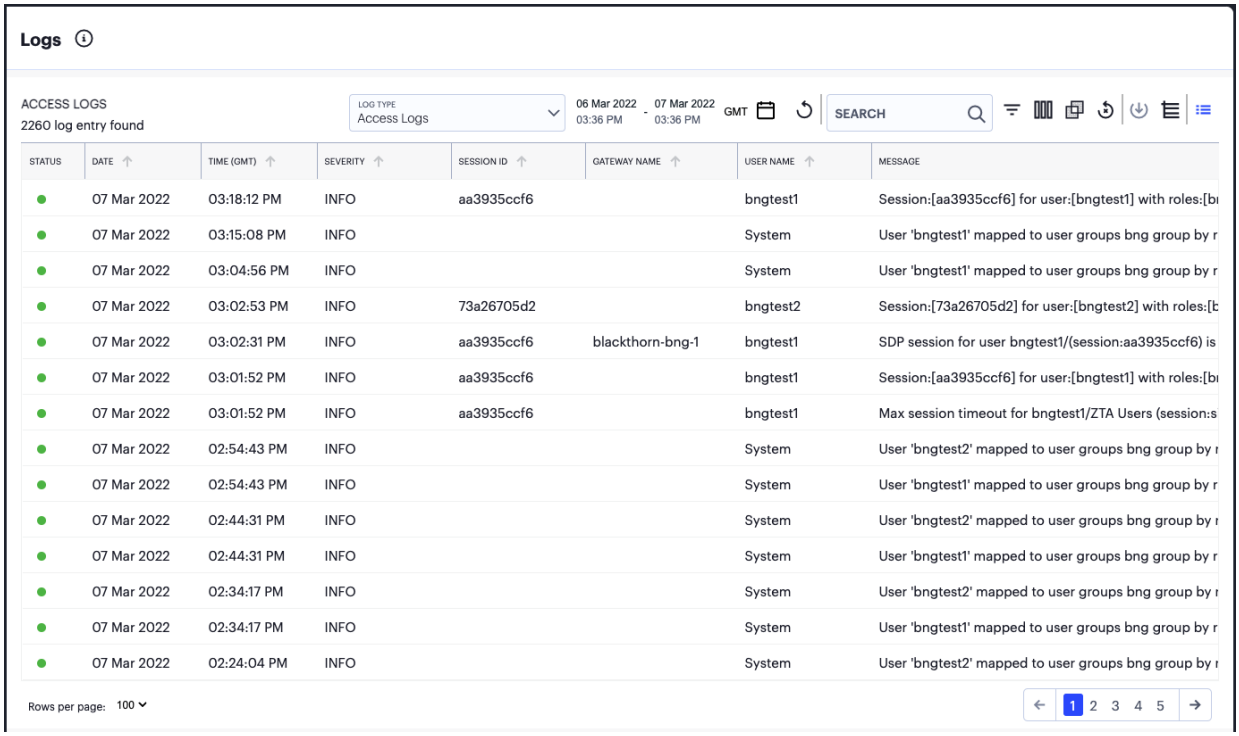

Viewing the Logs

This page comprises the following sections:

- The time period selector, see Setting a Log Time Period.
- Log selection and filtering controls, see Setting Log Criteria and Filtering the Output.
- The log record display, see Viewing Log Records.

> *nZTA* additionally provides a separate log records page pertaining to activity for specific *nZTA Gateway*s. To learn more, see Viewing and Monitoring Gateways in the Controller.

Setting a Log Time Period

Use the time period selector to set a time period or time range for your log results. Click the date-time display (highlighted) to show the selector dialog:

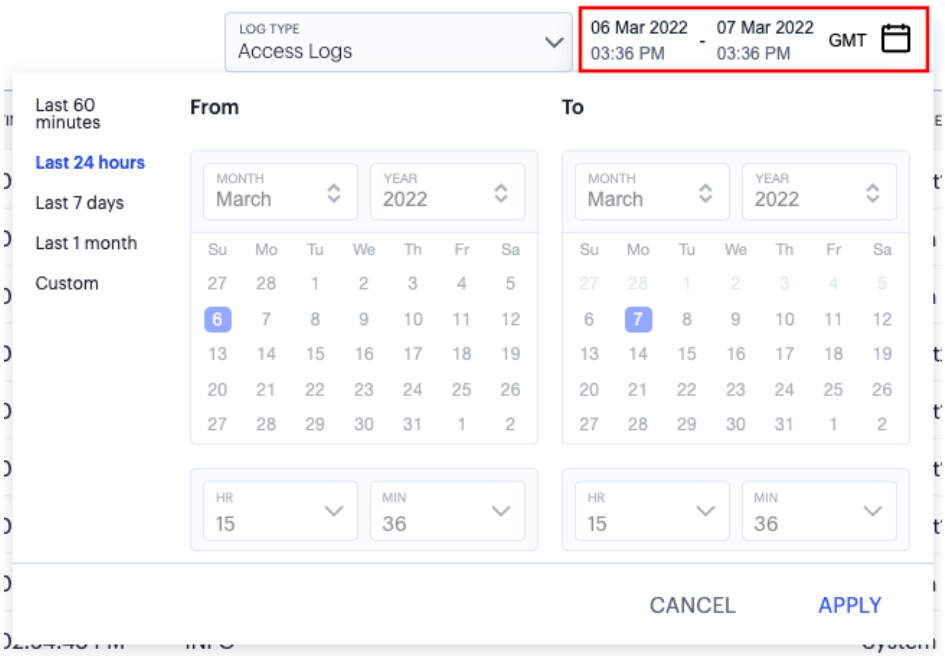

Setting a log time period

Set the time period you want to view using the available ranges at the top-left. Choose from:

- Last 60 minutes
- Last 24 hours (default)
- Last 7 days
- Last 1 month
- Custom

For **Custom**, set a specific *From* and *To* to denote the start and end of your custom date/time range.

> The date/time calendar controls are enabled for only the **Custom** option. However, the calendar continues to identify the applicable start and end date-time for all predefined time periods.

To apply your changes, click **Apply**. The selected time period is displayed in the filter bar and data on the page updates accordingly.

> To configure the timezone, see Setting the Timezone.

Setting Log Criteria and Filtering the Output

To set the criteria you want to use for viewing log data, use the controls above the main log display. This section also contains functions to highlight search terms, apply filters, and schedule log export jobs.

Select the primary log type you want to display by using the **Log Type** drop-down list:

Selecting a log type

Choose from:

- Access Logs
- Admin Logs
- Event Logs

Then, use the icons adjacent to the log selector to further control your log selection. Choose from the following:

- Logs are refreshed automatically by changing the criteria. To manually refresh the log display, click the following icon:

  

  Page refresh

- To search for a term in the displayed logs, click the following field:

  

  Search term highlighting

  nZTA highlights all matches in the log display.

- To trigger the advanced filter selection, use the following icon:

  

  Advanced Filtering

  To learn more, see Filtering the Logs.

- To change the fields displayed for each log line, click the following icon:

  

  Show or hide log fields

  In the field selector, click a field name to toggle between show or hide. A *tick* icon indicates a displayed field. After you are finished, click the context menu icon to close the selector. See Viewing Log Records.

- To apply grouping to the displayed log records, click the following icon:

  

  Group log records by selected criteria

  This feature applies grouping to a selected field in the log record display, such that records are accumulated and grouped together under each unique data item identified in that field. Through grouping, an admin can quickly view the number of records of a particular type.

  To learn more about record grouping, see Viewing Detailed Logs for a Chart.

- To remove any applied filters from the data set, click the following icon:

  

  Remove any applied filters from the data

- To export the displayed log as a CSV or JSON text file, or to set up a new scheduled log export job, click the following icon:

  

  Export filtered logs

  To learn more about log export jobs, see Exporting Logs.

- To view the status of currently-scheduled log export jobs, click the following icon:

  

  View scheduled log export jobs

- To learn more about log export jobs, see Exporting Logs.

- To change the view density, click the following icon:

Switching between default and dense log record views

Viewing Log Records

The main part of the page shows the log records that match your selected criteria. The number of matching log records is displayed at the top-left.

Each log line includes the following fields:

- A status indicator showing the level of severity associated with each log event. Use the following table for a guide to the meaning of each indicator color:

Severity	Status Color
INFO	Green
MINOR	Amber
MAJOR	Amber
CRITICAL	Red

- The date and time of the event.
- The message ID that identifies this type of event.
- The severity of the event in words.
- The session ID that was the source of the event, where applicable.
- The ID of the *nZTA Gateway* that reported the event, where applicable.
- The name of the *nZTA Gateway* that reported the event, where applicable.
- The IP address identified as the source of the event.
- The user name associated with the event, where applicable.
- The ID of the device associated with the event, where applicable.
- The message (description) of the event.

Use the page controls at the bottom of the window to select the number of log records/rows per page:

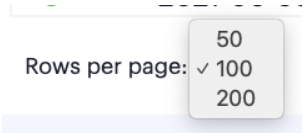

Setting the number of log rows per page

Choose from:

- 50
- 100 (default)
- 200

To cycle through the log pages, use the page controls at the bottom-right.

Where a single log message is too long for the display, use your pointing device to scroll the optional fields display to the left or right.

Furthermore, to view a single log entry in a dedicated panel, click the log message text to activate the info-panel view:

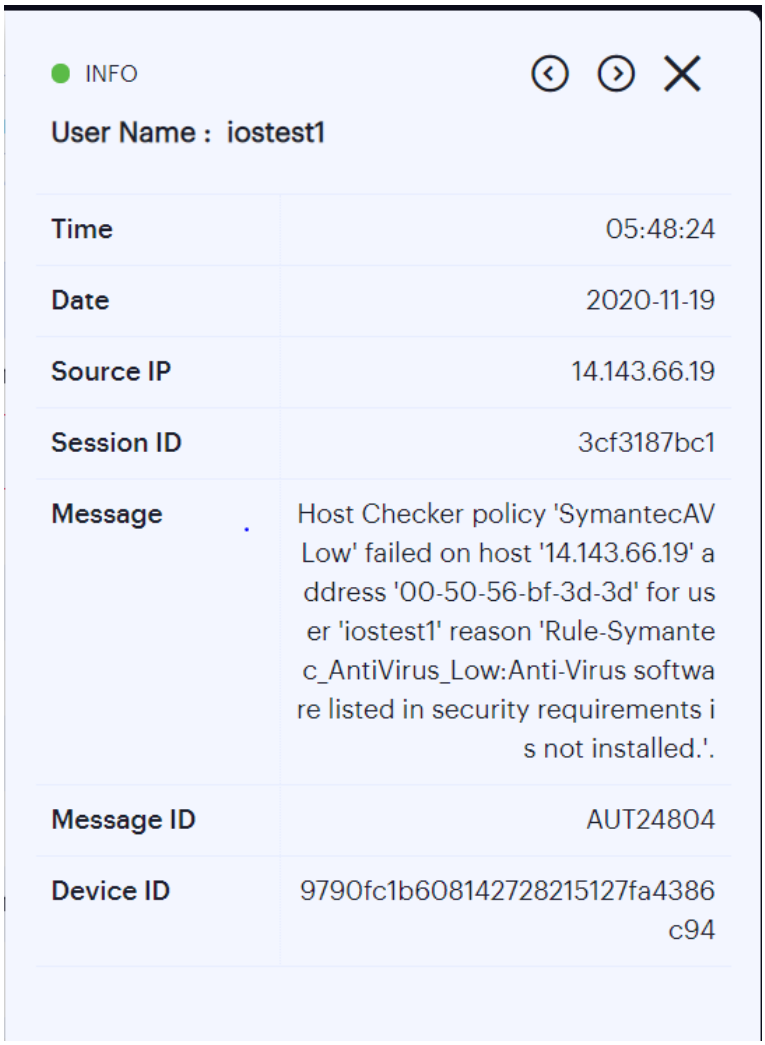

Viewing a single log entry in the info-panel

In the info-panel, use the **Previous** and **Next** icons to cycle through each log entry in turn.

Filtering the Logs

The **Logs** page provides an advanced field filter through which you can narrow down the displayed log entries to a sub-set that matches the filters you apply. You can also save filter definitions for later use.

To set a filter, click the following icon:

Activating the advanced filter

Next, use the side-panel dialog to add one or more new field filters.

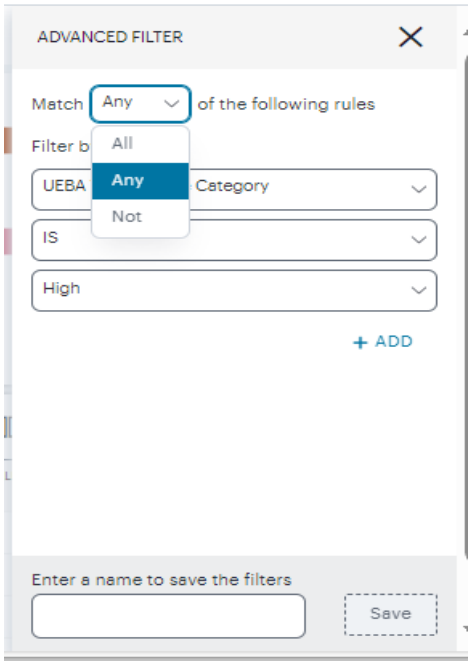

Adding a new log filter

In this dialog, select a matching criteria for the filters.

- **All**: performs AND operation on the filters
- **Any**: performs OR operation on the filters
- **Not**: Negates the list of filters

You can recall a saved filter through the **Saved Filters** drop-down list or set new filter criteria through the **Filter by** section.

By selecting a saved filter, the filter criteria are populated into the panel. To then apply the saved filter, click **APPLY FILTER**.

You can add additional criteria lines to a recalled filter before applying it, but the saved filter is unaffected.

Saved filters are preserved across all log pages in the Tenant Admin portal, but might not be valid for all pages. For example, a saved filter created on the **Insights > Logs** page might not be applicable to the data on the **Gateways > Logs** page (in other words, where a filter references a log field not applicable to *nZTA Gateways*). In this case, where you attempt to select an invalid filter, *nZTA* presents an error.

When setting new filter criteria, use the **Selector** drop-down list to choose the field you want to filter on, add an **Operator** type, and then enter the **Value** you want to apply. For the operator, choose from:

- **IS**: The selected field matches exactly the value you specify.
- **CONTAINS**: (where applicable) The selected field contains as a sub-string the value you specify.

To add further criteria to this filter, click the *plus* symbol. Then, repeat the above step as desired. To remove a criteria line, click the corresponding *X* icon.

To apply the defined filter, click **APPLY FILTER**.

Your filters remain in place through data refreshes, and active filters are identified by the *Filters are applied on this page* label at the top of the page. To remove a filter, click the filter icon (or the link at the top of the page) to re-display the filters side-panel dialog. Then, click **CLEAR ALL** to remove all active filters.

To save a filter for future use, use the save-as facility at the bottom of the panel. Enter a name for your saved filter in the text box provided, then click **Save**. You can recall your filter through the *Saved Filters* list at the top of the panel.

To delete saved filters, use the *Saved Filters* list. Select the check box adjacent to the filter, or filters, you want to delete, then click **DELETE** from the bottom of the panel.

Exporting Logs

nZTA provides the ability to export the currently-displayed log as a Comma-Separated Value (CSV) or JavaScript Object Notation (JSON) text file. You can download the log immediately or set up a scheduled job to activate or repeat the export action at a defined time and interval of your choosing.

To access the Export Logs page:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Insights** icon, then select **Logs**.

   The *Logs* page appears.

3. Select the log type you want to display in the **Log Type** drop-down list. Choose from:
    - Access Logs
    - Admin Logs
    - Event Logs
4. Click the *cloud* icon at the top of the page:

    

    Accessing the Export Logs Settings page

    The Export Logs page appears:

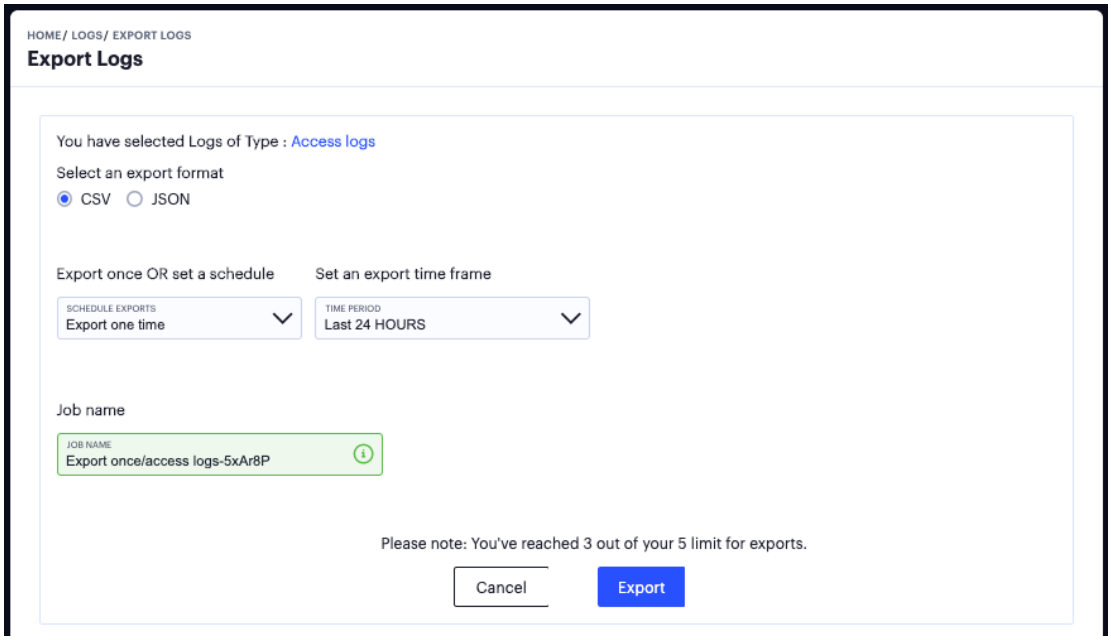

The Export Logs settings page

Use the Export Logs settings page to configure an export operation, either to execute immediately as a one-off job, or as a scheduled job.

Configure the following settings:

- Select either **CSV** or **JSON** as the output format.
- Select the frequency of the export operation. Choose from:
    - **Export one time**: Perform the log export now as a single job.
    - **Daily data export**: Create a daily export job executed once per day from the selected start date, up to and including the stop date (if defined).
    - **Weekly data export**: Create a weekly export job executed once per week on the selected start day, up to and including the stop date (if defined).
    - **Monthly data export**: Create a monthly export job executed once per month on the selected start day, up to and including the stop date (if defined).

  If a stop date is specified, this is the date the schedule ceases. In the case of weekly or monthly jobs, if this date falls before the expected run date for that period, the job is terminated without running. For example, in a weekly run scheduled to execute every Thursday, if the stop date is set as a Tuesday, the final run of the job would be the previous Thursday.

  

  A daily data export job continues to run for one extra day beyond the selected end date in order to process the logs for the final scheduled day.

  

  For daily/weekly/monthly frequency export jobs, *nZTA* allows for a maximum of 5 runs per scheduled export job. That is, each schedule runs a maximum of 5 times. On the sixth run, the first run is deleted (together with the log file), and so on.

- Set an export time frame. For one-time exports, choose from:
    - **Last 60 minutes**
    - **Last 24 hours**
    - **Last 7 days**
    - **Last 1 month**
    - **Set a date range (30d max)**: This option presents a configurable start and end date.

  For daily, weekly, and monthly exports, this option switches to show start and end date parameters. You do not need to specify an end date; in this case, the job remains active until deleted.

- Enter a **Job name** for the export operation. *nZTA* suggests an appropriate name; use this, or type your own.
- To execute the defined job, click **Export**.

To view all scheduled export logs jobs, and to download the log files created by each job, see Viewing Scheduled Log Export Jobs and Downloading Log Files.

> *nZTA* allows for a maximum of 5 defined export jobs. Each job that you add reduces the total, as displayed at the bottom of the page. This is a separate limit to the maximum number of job runs described earlier.

Viewing Scheduled Log Export Jobs and Downloading Log Files

To view the status of your current log export jobs:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nZTA* menu, click the **Insights** icon, then select **Logs**.

   The *Logs* page appears.
3. Click the *list* icon at the top of the page:

   

   Accessing the Job Status page

   The Job Status page appears:

   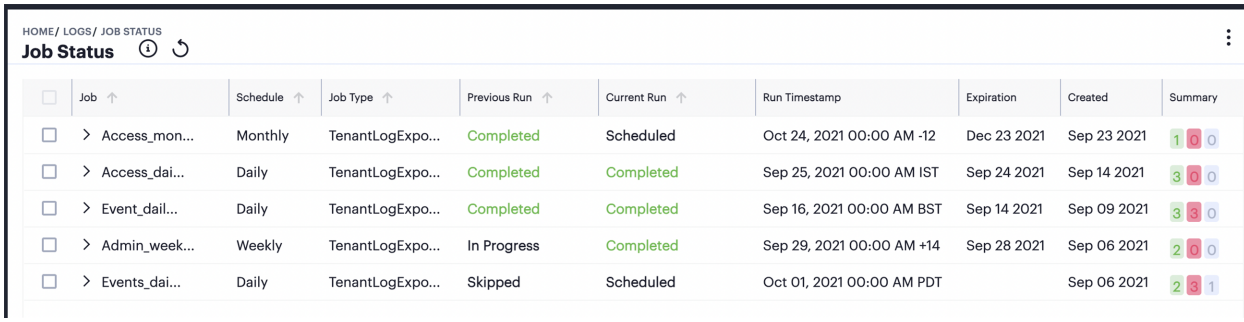

Job	Schedule	Job Type	Previous Run	Current Run	Run Timestamp	Expiration	Created	Summary
Access_mon...	Monthly	TenantLogExpo...	Completed	Scheduled	Oct 24, 2021 00:00 AM -12	Dec 23 2021	Sep 23 2021	1 0 0
Access_dai...	Daily	TenantLogExpo...	Completed	Completed	Sep 25, 2021 00:00 AM IST	Sep 24 2021	Sep 14 2021	3 0 0
Event_dail...	Daily	TenantLogExpo...	Completed	Completed	Sep 16, 2021 00:00 AM BST	Sep 14 2021	Sep 09 2021	3 3 0
Admin_week...	Weekly	TenantLogExpo...	In Progress	Completed	Sep 29, 2021 00:00 AM +14	Sep 28 2021	Sep 06 2021	2 0 0
Events_dai...	Daily	TenantLogExpo...	Skipped	Scheduled	Oct 01, 2021 00:00 AM PDT		Sep 06 2021	2 3 1

   The Job Status page

Use the Job Status page to:

- View the status and progress of currently scheduled log export jobs.
- Download log files for completed job runs.

For each job on the Job Status page, you can view the configured details of the export operation along with status indicators for progress of the previous and outstanding job runs.

> A job run refers to a single run of a scheduled job. For example, in a weekly data export job, a job run refers to the export operation scheduled or completed for one specific week within the start and end dates. Thus, a scheduled log export job is comprised of one or more job runs.

The **Summary** column provides totals of successful job runs, unsuccessful/failed job runs, and inactive job runs.

Click any of the fields in a single job row to display an info-panel at the side showing more details about the scheduled job:

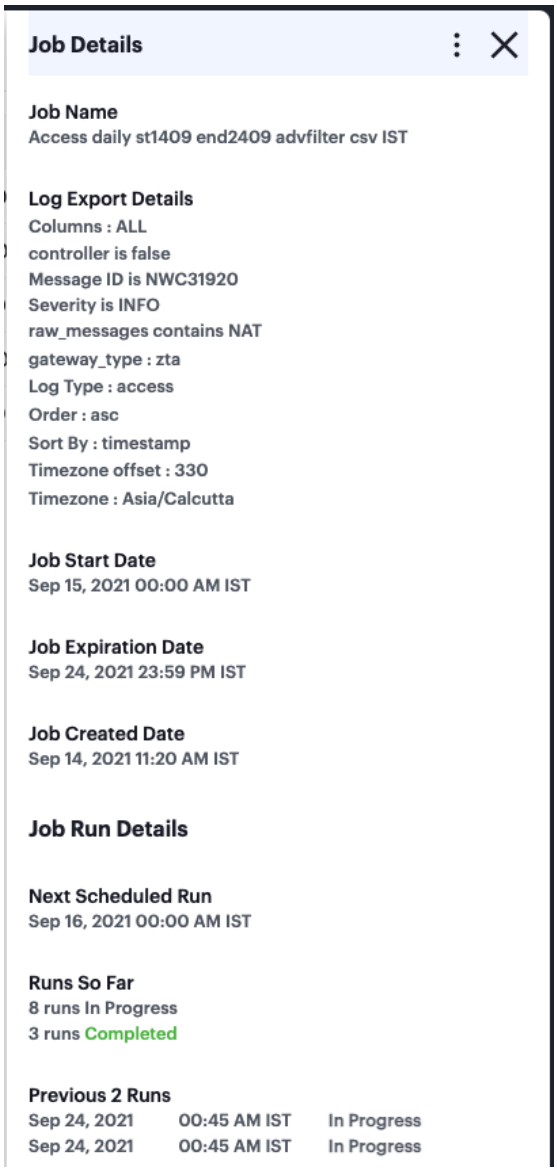

The Job Details info-panel

To access the log files and view more information about each individual job run, click the down-arrow adjacent to the *Job* name:

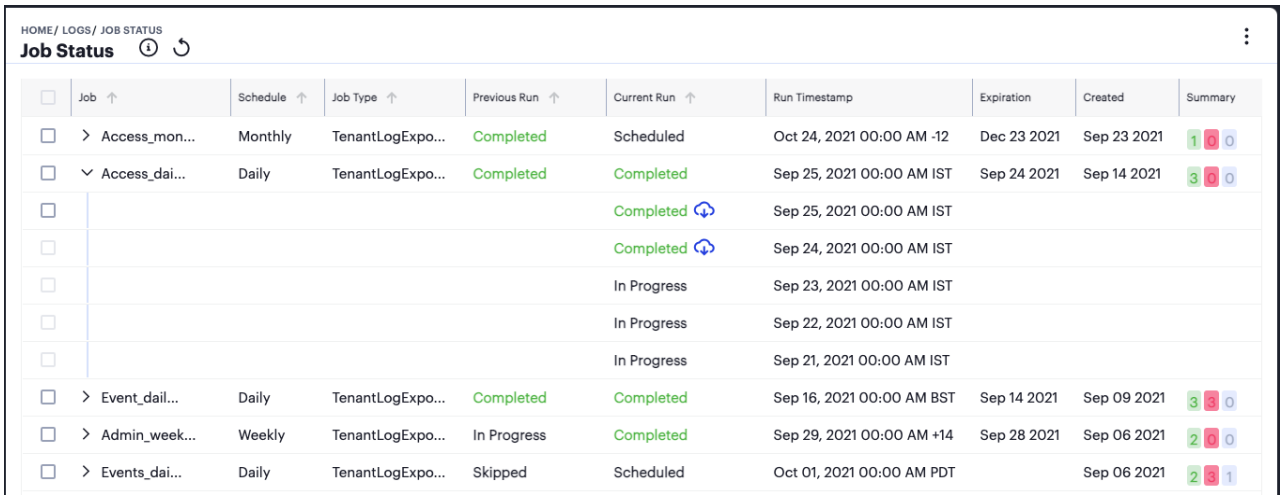

Job	Schedule	Job Type	Previous Run	Current Run	Run Timestamp	Expiration	Created	Summary
Access_mon...	Monthly	TenantLogExpo...	Completed	Scheduled	Oct 24, 2021 00:00 AM -12	Dec 23 2021	Sep 23 2021	1 0 0
Access_dai...	Daily	TenantLogExpo...	Completed	Completed	Sep 25, 2021 00:00 AM IST	Sep 24 2021	Sep 14 2021	3 0 0
				Completed	Sep 25, 2021 00:00 AM IST			
				Completed	Sep 24, 2021 00:00 AM IST			
				In Progress	Sep 23, 2021 00:00 AM IST			
				In Progress	Sep 22, 2021 00:00 AM IST			
				In Progress	Sep 21, 2021 00:00 AM IST			
Event_dail...	Daily	TenantLogExpo...	Completed	Completed	Sep 16, 2021 00:00 AM BST	Sep 14 2021	Sep 09 2021	3 3 0
Admin_week...	Weekly	TenantLogExpo...	In Progress	Completed	Sep 29, 2021 00:00 AM +14	Sep 28 2021	Sep 06 2021	2 0 0
Events_dai...	Daily	TenantLogExpo...	Skipped	Scheduled	Oct 01, 2021 00:00 AM PDT		Sep 06 2021	2 3 1

Showing all job runs for a scheduled export job.

> For daily/weekly/monthly frequency export jobs, *nZTA* allows for a maximum of 5 runs per scheduled export job. That is, each schedule runs a maximum of 5 times. On the sixth run, the first run is deleted (together with the log file), and so on.

As with a scheduled job, click on any of the fields in the job run row to display an info-panel at the side showing more details about the job run:

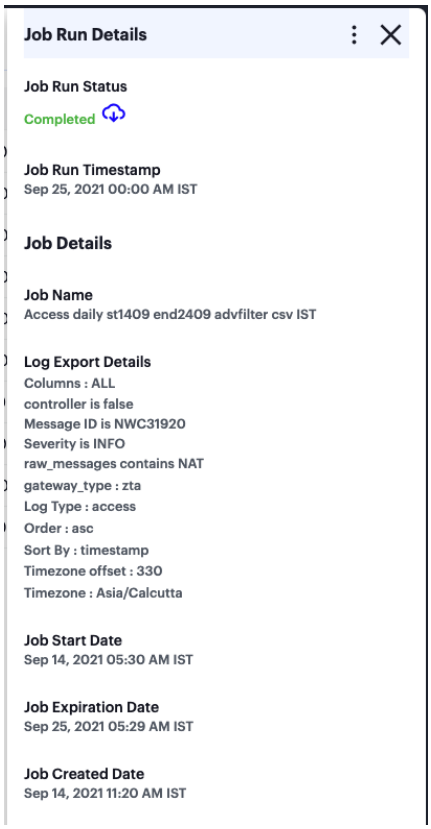

The Job Run Details info-panel

To download the log file generated by the job run, click the *cloud* icon for a completed job run:

Downloading a log file

To remove a scheduled log export job, or any of the completed job runs within the job, tick the checkbox adjacent to the job/job run and then access the context menu at the top of the page:

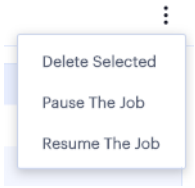

The Job Status menu

Select from the following options:

- **Delete Selected**: Remove all jobs or job runs that have been selected.
- **Pause the Job**: Instruct the outstanding job runs in the schedule to become inactive. The schedule continues chronologically, but no further log export operations are completed while in this state.
- **Resume the Job**: Resume the schedule starting at the next scheduled job run.

If you choose to delete a complete job, all job runs and log download files are removed permanently.

Associating Geographical locations to IP Addresses

nZTA provides the mapping of Gateway geographic location to IP address.

Before you start, make sure that you have the following information:

- The public IP address/range for the Gateway. This is the IP address at which clients can externally reach the Gateway.
- The Gateway geographic location information such as country, state/province and city.

To add a new location:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.

   The *Network Overview* page appears by default.

2. From the *nZTA* menu, click the **Administration** icon, then select **Custom Geo IP**.

   The *Custom Geo IP* page appears. This page lists all defined geographical associations to IP addresses.

1. Click "+" at the top of the page:

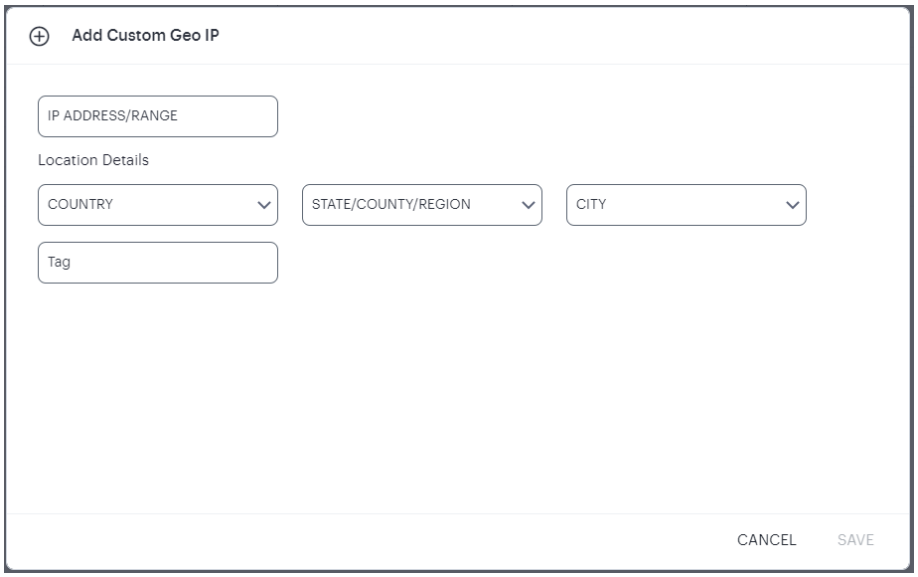

The Custom Geo IP

2. Enter the **IP Address/range**.
3. Select the **Country**.
4. Select the **State/Province**.
5. Select the **City**.
6. Enter a **Tag** for this IP Address/range.
7. Click **Save**.

Actions

nZTA enables you to configure actionable insights, such that when certain conditions are met a defined action is executed.

To configure an action:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.

   The *Network Overview* page appears by default.

2. From the *nZTA* menu, select the **Insights** icon, then select **Actions**.

   The *Actionable Insights* page appears.

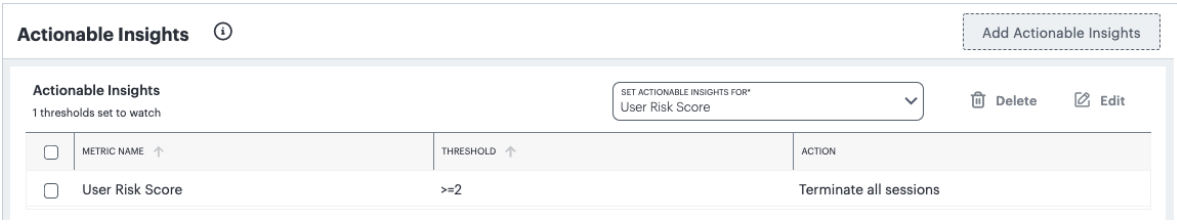

Viewing Actionable Insights

Use this page to view and configure actions that are triggered by a condition being met.

The following conditions are supported in this release:

- **UEBA Threat Score**: If a user's UEBA Threat score breaches a set threshold, the selected action is triggered.

> The condition remains in force until the user's UEBA Threat score is manually reset. To learn more about resetting a UEBA Threat score, see Viewing a Summary of UEBA Threat Scores for your Users.

The following actions are supported in this release:

- **Terminate all existing sessions for the user**: If the set condition is reached, all sessions for the affected user are terminated. If that user attempts a further login, *Ivanti Secure Access Client* denies the attempt and displays a message concerning the breach, directing the user to contact their administrator. *nZTA* also records an admin log event referencing the fact (see Checking the Logs).

To add a new condition:

1. Select **Add Actionable Insight**.
2. In **Set Actionable Insight for**, select a condition to apply.

   The configurable options for that condition are displayed.

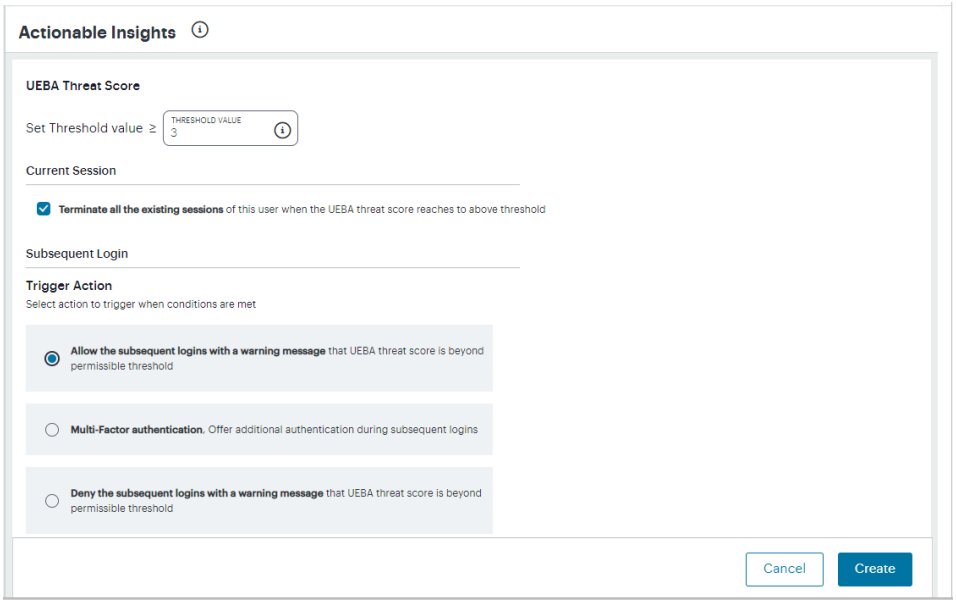

Add Actionable Insights

3. Set the required options/thresholds for the condition.
4. In **Trigger Action**, select the applicable action to be applied if the condition is met.
5. From the Subsequent Login section, select one of the following actions to trigger when conditions are met:
   - Allow subsequent logins with a warning message
   - Offer Multi-factor Authentication during the subsequent logins
   - Deny subsequent logins with a warning message
6. To save your changes, select **Create**.

To edit or delete an actionable insight, select the check box adjacent to the desired condition and select **Delete** or **Edit** as applicable.

When the user sessions are terminated due to reaching the threshold UEBA Threat score, the admin log messages are generated in nSA. Select the **Logs** tab to view the list of log messages.

Reports

nZTA provides the ability to generate and download activity reports from pre-defined report templates or through a custom defined report. It also supports scheduling the reports, to be generated either daily, or weekly once, twice or thrice.

To access the **Reports** page:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.

   The *Network Overview* page appears by default.

2. From the *nZTA* menu, click the **Insights** icon, then select **Reports**.

   The *Reports* page appears, on the *Report Templates* tab.

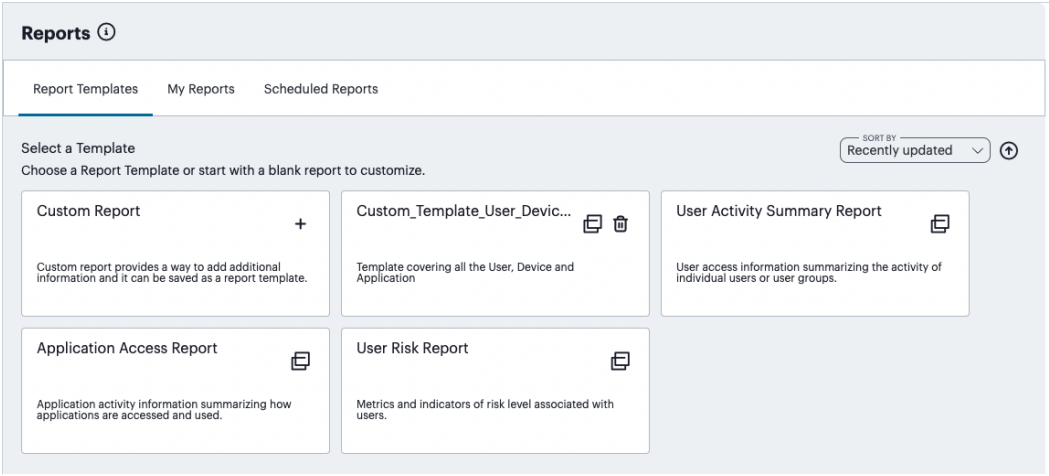

Viewing report templates

The *Reports* page provides the following tabs:

- **Report Templates**: Contains all built-in and custom-saved templates upon which all reports are based, including a *Custom Report* option to enable creation of customized reports.

  Built-in templates are read-only whereas custom-saved templates added by a tenant admin can be deleted. You can identify custom templates as per the following image:

  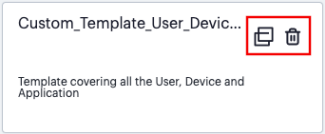

  Identifying a custom-created report template from a built-in template (the delete option - indicated)

- **My Reports**: Contains all generated report instances:

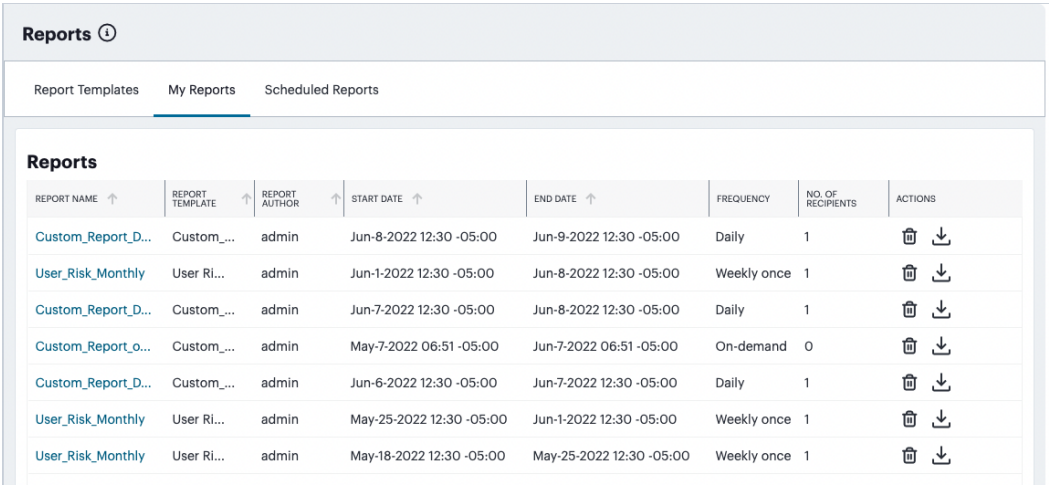

The list of generated reports

Reports shown on this page either originate from a singular on-demand request, or represent an instance of a scheduled report run. For example, if you schedule a report to run daily, at the requisite time each day a new instance of the report is generated and placed here.

For each generated report, you can:

- Select the report name to view a summary of the configured parameters:

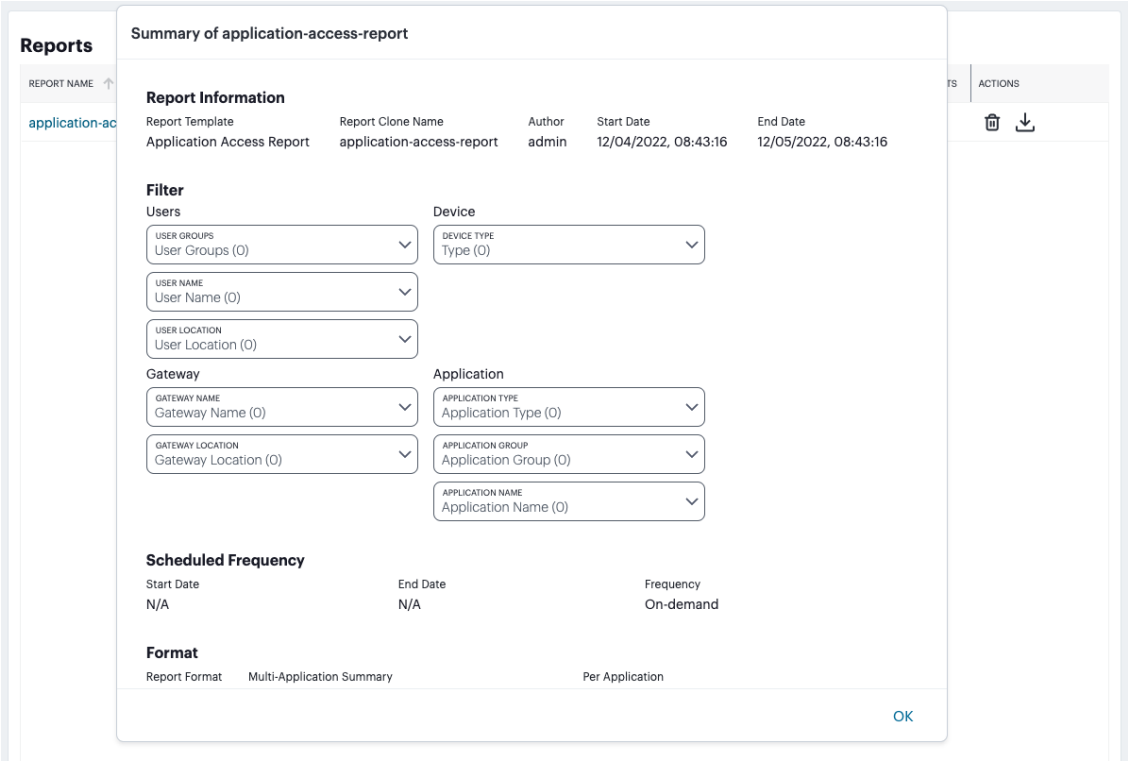

Report parameters

- In the *Actions* column, select the download icon to view and download the report in the specified format (PDF, JSON, or CSV)

The Download icon

- In the *Actions* column, select the delete icon to permanently remove the report instance.

The Delete icon

- **Scheduled Reports**: Contains the list of report schedules:

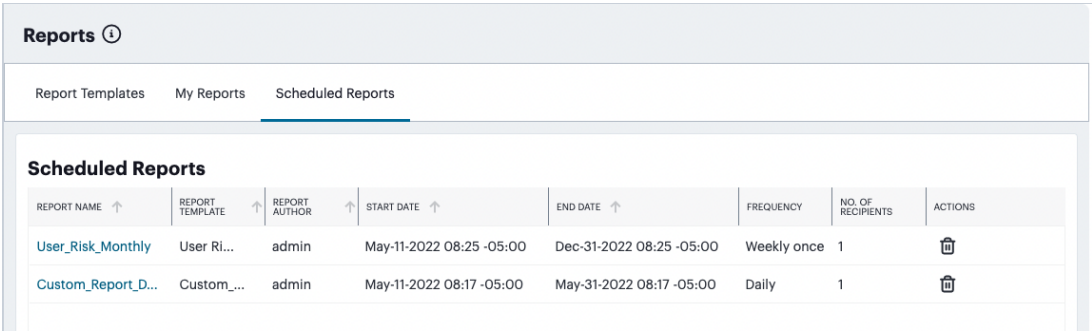

The list of report schedules

Each entry on this page represents a scheduled report definition. For each entry, you can:

- Select the report name to view a summary of the configured parameters:

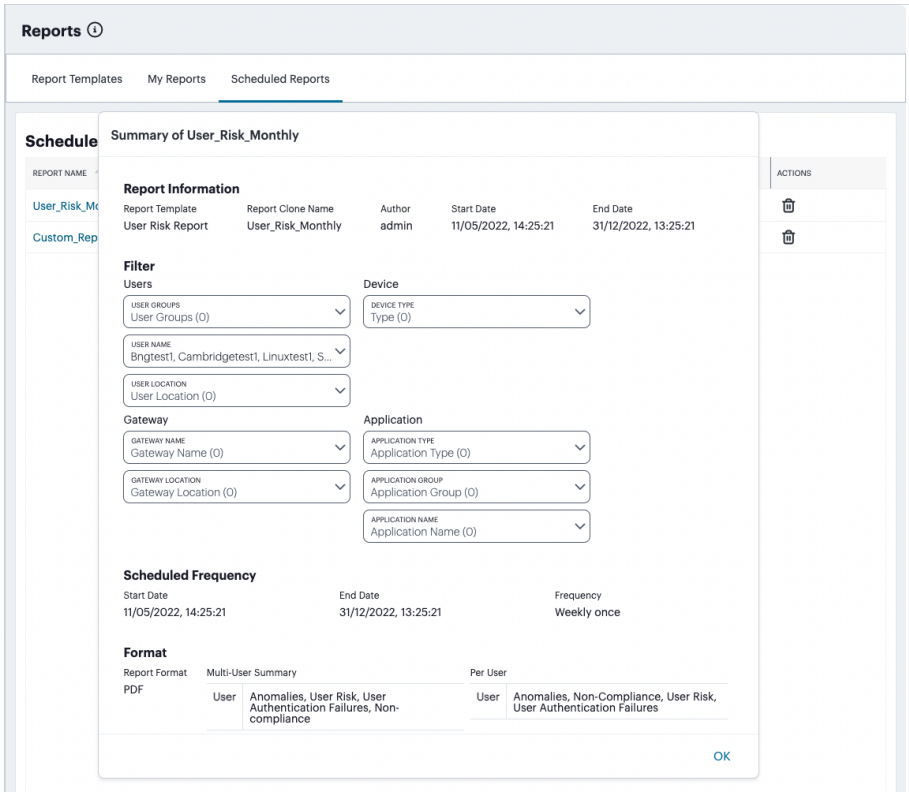

Schedule details

- In the *Actions* column, select the delete icon to remove the schedule:

The Delete icon

Creating a Report

This section describes how to create a new report. You can choose to create the report based on one of the following methods:

- Create a new custom report
- Create a new report based on one of the built-in predefined report templates provided as a part of your subscription
- Create a new report based on a custom template created by a tenant admin

To configure a report:

1. On the **Report Templates** tab, choose the template option from which to create your report.

   To add a new custom report, select the "Custom Report" option:

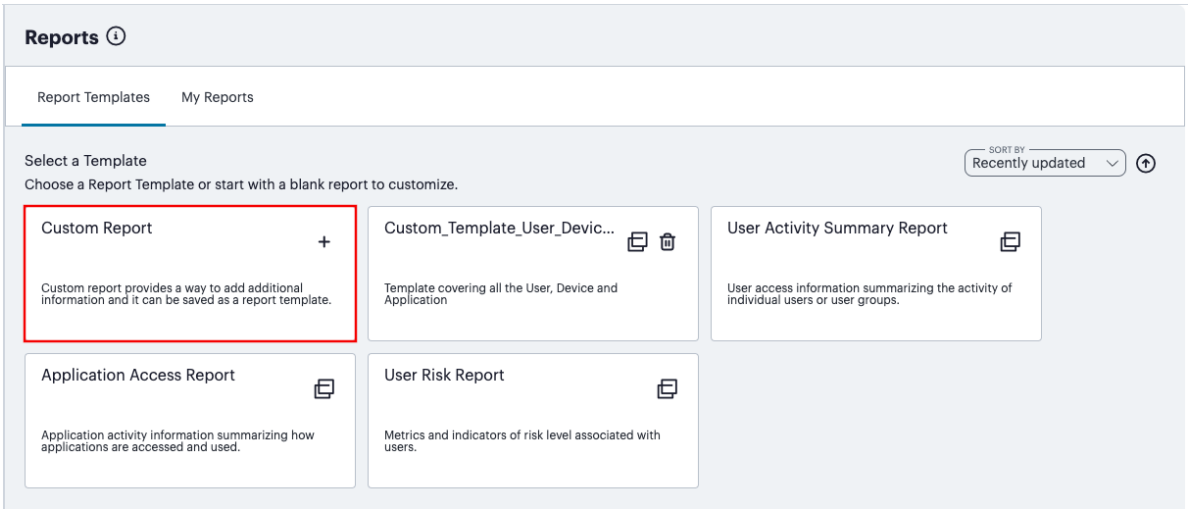

   Adding a new custom report

   To add a report based on a template, select the built-in or custom template of your choice:

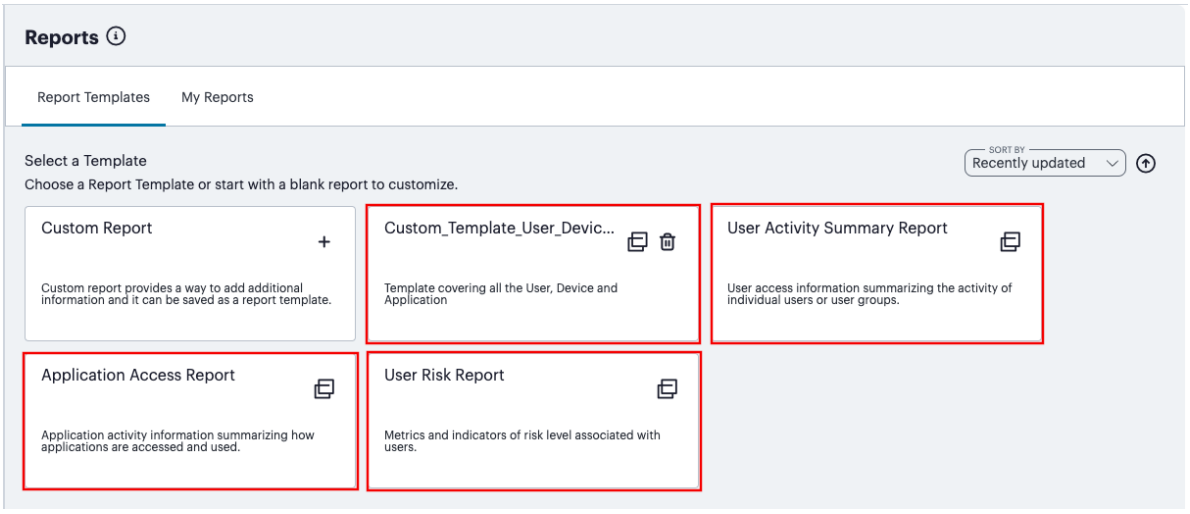

   Adding a new report based on a template

   The report wizard appears, beginning with the **Clone** step:

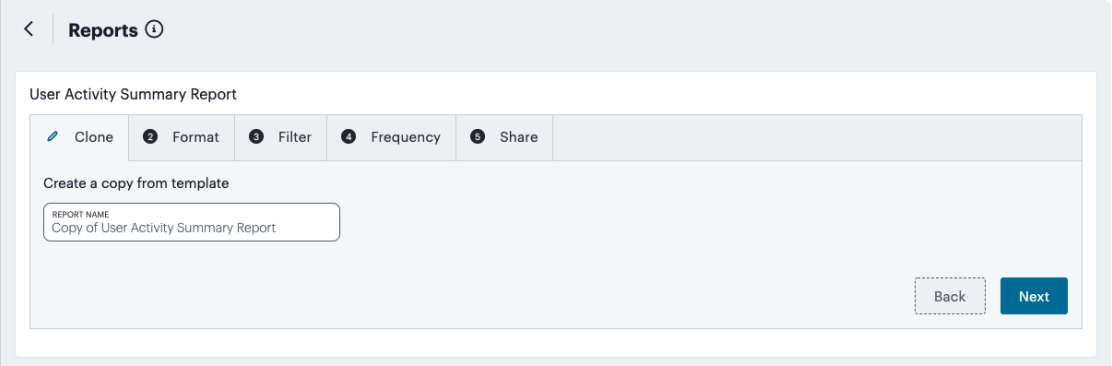

Creating a report - Clone step

2. Enter a unique name for the report and click **Next** to continue.

3. In the **Format** step:

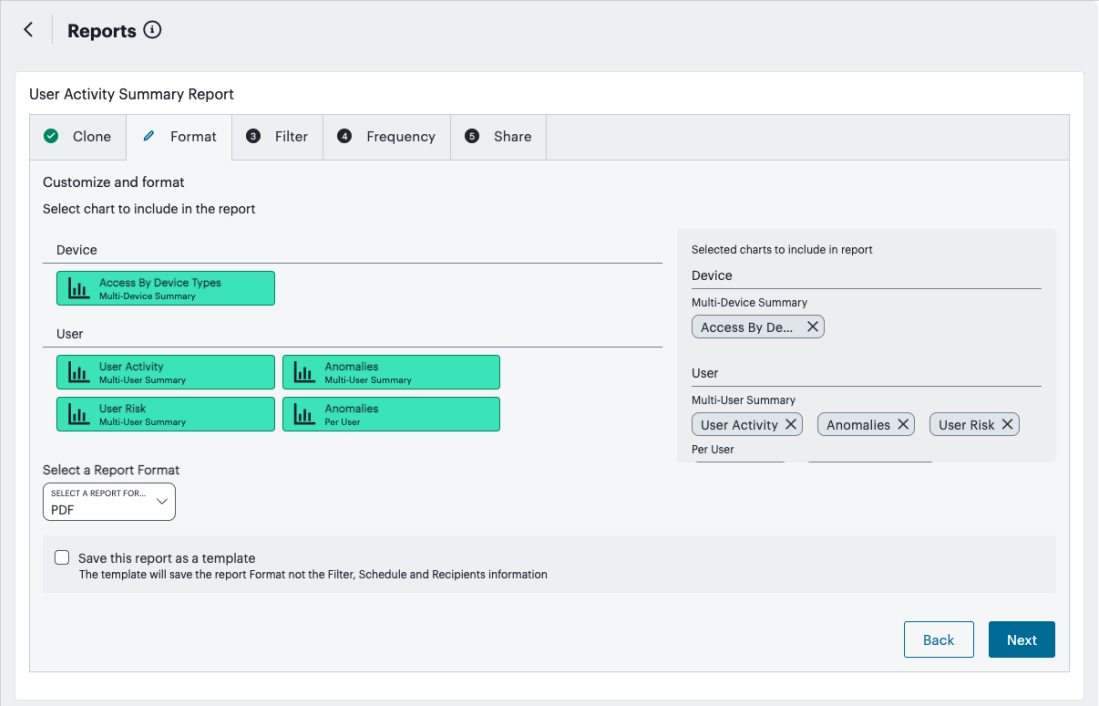

Creating a report - Format step

- Select or deselect the required charts from the *User*, *Device*, and *Application* sections as applicable. Selected items appear in the right-hand panel.

  Use your pointing device to vertically scroll the charts panel as required.

- Select the report format (PDF, JSON, or CSV).
- (Optional) select **Save this report as a template** to create a new custom template containing your selections. Enter a template name and description in the fields provided.

Click **Next** to continue.

4. In the **Filter** step, for each category of Users, Devices, Gateways, and Applications, select or deselect the named items you want to include. For example, within *Users*, use the drop-down controls to select specific *User Groups*, *User Names*, or *User Locations* you want to include in the report:

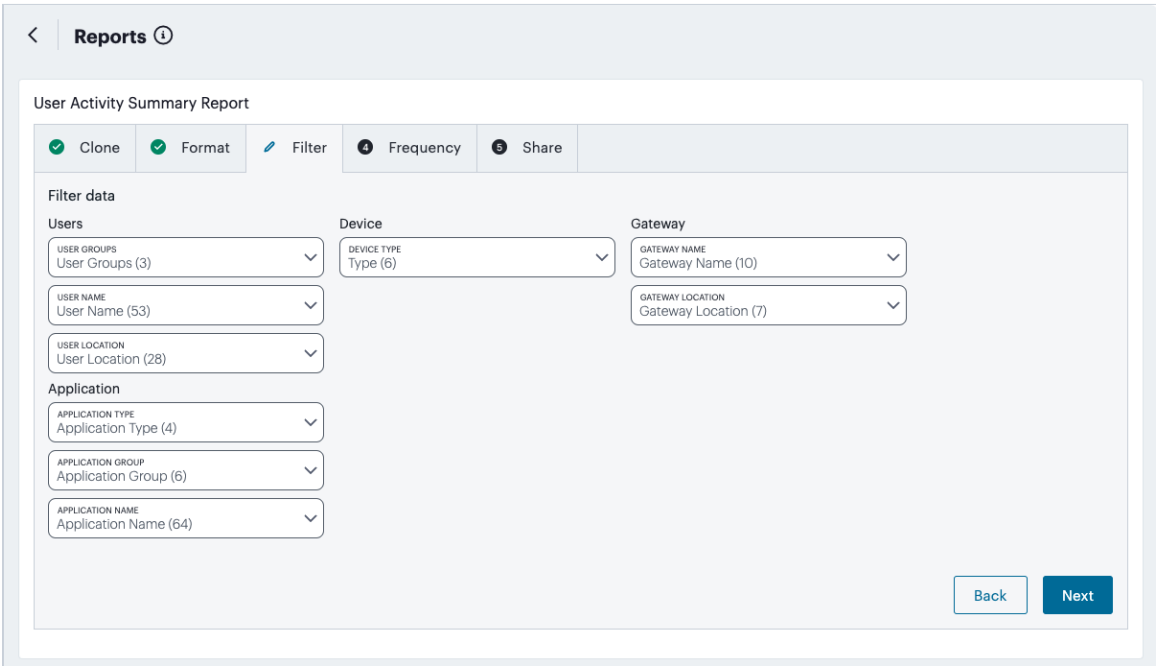

Creating a report - Filter step

Objects that appear in the drop-down lists in this step are derived from those items last accessed within the previous 30 days only. Items last accessed earlier than this time are not shown.

Click **Next** to continue.

5. In the **Frequency** step, set the frequency with which you want this report to run:

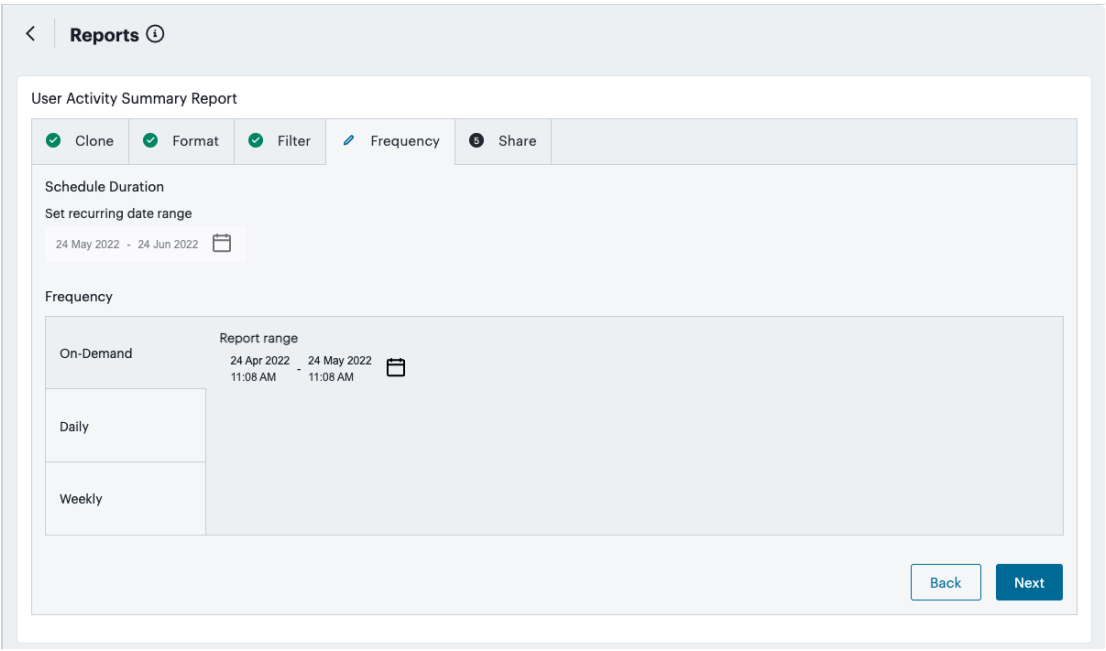

Creating a report - Frequency step

Choose from:

- **On Demand**: Run once for a specified date and time period
- **Daily**: Run daily at a defined time
- **Weekly**: Run at a specified time on certain days of the week

For Daily and Weekly, use **Set recurring date range** to set the start and end dates for which you want the schedule to run.

Click **Next** to continue.

6. In the **Share** step, add the recipients with which the report should be shared (if applicable):

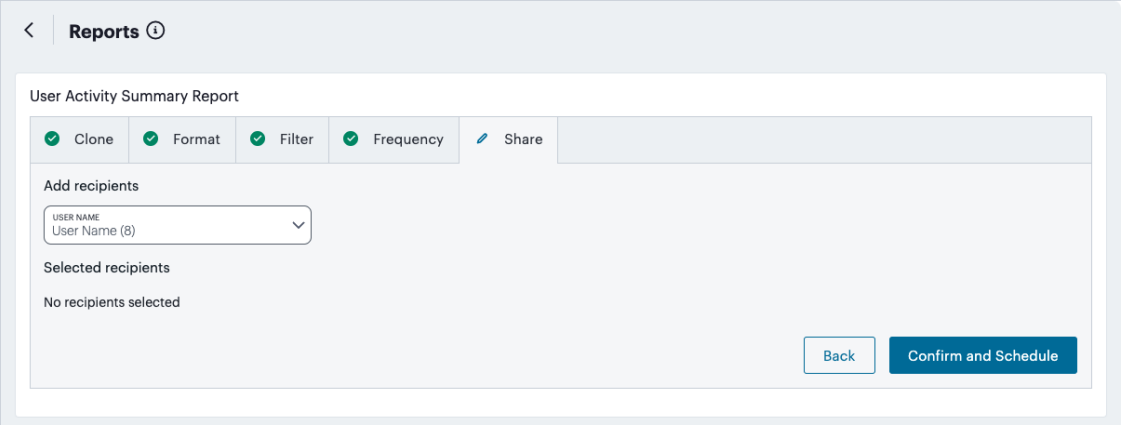

Creating a report - Share step

7. To complete the wizard and schedule the report according to the selections made, select **Confirm and Schedule**.

If you elected to save the report as a custom template during the **Format** step, the new template is displayed on the **Report Templates** tab.

Viewing Alerts and Notifications

The **Alerts** page lists all alerts and notifications that have been raised by *nZTA*.

To view the **Alerts** page, click the **Alerts** icon and then click **See all Alerts**:

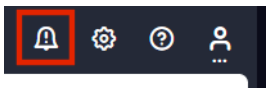

Alerts icon

The **Alerts** page appears. For example:

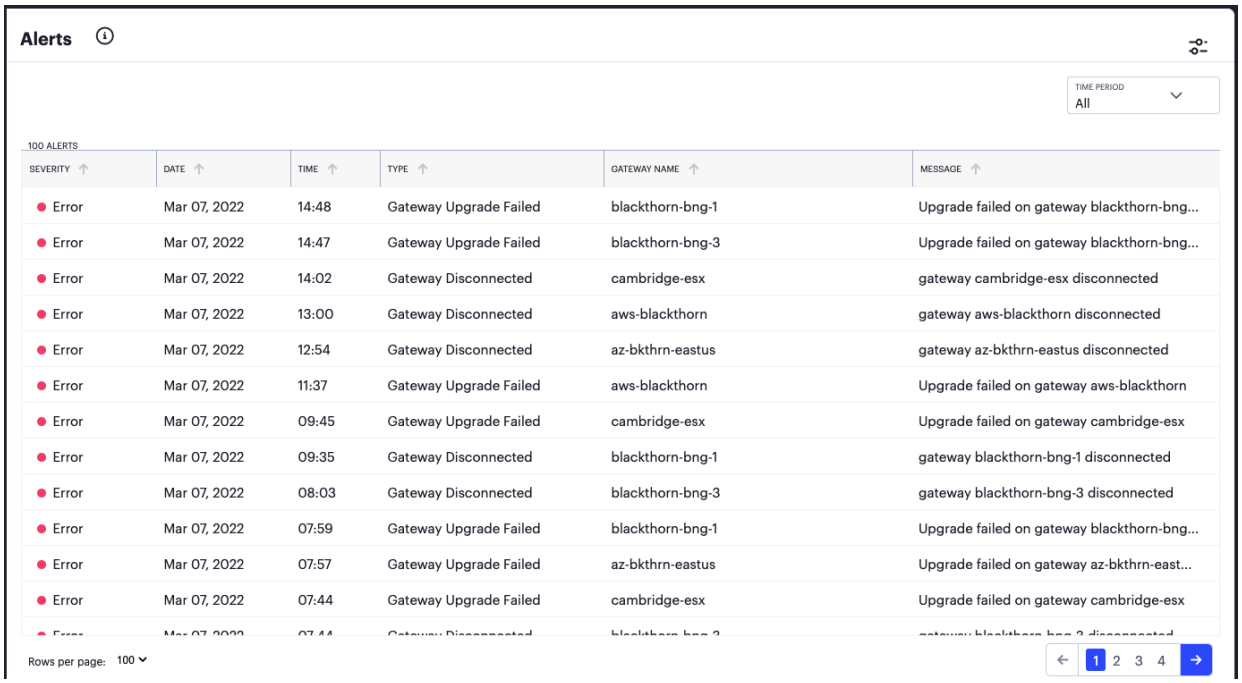

Alerts page

The alerts table supports the following alert types:

- AAA Config Pull Failure
- AAA Config Pull Success
- AAA Config Pull Success - Failure Resolved
- AAA Journal Update Failed
- AAA Journal Update Success
- Config Sync Rule Deleted
- Config Sync Rule Updated
- Config Sync Target Cluster Deleted
- Custom Domain Certificate for mTLS Domain Due for Renewal
- Custom Domain Certificate for mTLS Domain Expired
- Custom Domain Certificate for TLS Domain Due for Renewal
- Custom Domain Certificate for TLS Domain Expired

- Device Vulnerability Risk Rating (VRR) Critical
- Device Vulnerability Risk Rating (VRR) High
- Device Vulnerability Risk Rating (VRR) Medium
- Device Vulnerability Risk Rating (VRR) Low
- Gateway Config Apply Failed
- Gateway Config Import Failed
- Gateway Disconnected
- Gateway Invalid Configurations Cleared
- Gateway Upgrade Failed

To filter the alerts table by type:

1. Click **Configure Alert Rules** icon.

   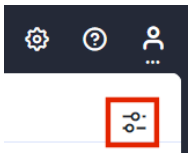

   Configure Alert Rules icon

   The **Configure Alerts & Notifications** page appears.

2. Click **Alert Types** and select the required type.
3. Click **Close**.

To filter the alerts table by time period, click **Time Period** and select the required time period.

To sort the alerts table into ascending or descending order of a specific property, click on one of the following column headings in the alerts table:

- **Severity**
- **Type**
- **Message Type**

Using Enterprise Integration to Export Your Logs for External Analysis

Introduction and Prerequisites

Ivanti Neurons for Zero Trust Access (*nZTA*) gathers data for events that occur on your *nZTA Gateways* and from the *Controller*'s Authentication, Authorization and Accounting (AAA) service. These events are categorized and stored in the following log files:

- Access Logs
- Admin Logs
- Event Logs

You can view and obtain these logs at any time through the *nZTATenant Admin Portal* (for details, see Checking the Logs).

nZTA enables you to automatically export log data for analysis in an external third-party enterprise analytics or visualizer tool through *Enterprise Integration*. This enables enhanced visibility into the health and efficiency of the services running in your *nZTA* deployment, or to facilitate debugging in the event of unexpected service behavior.

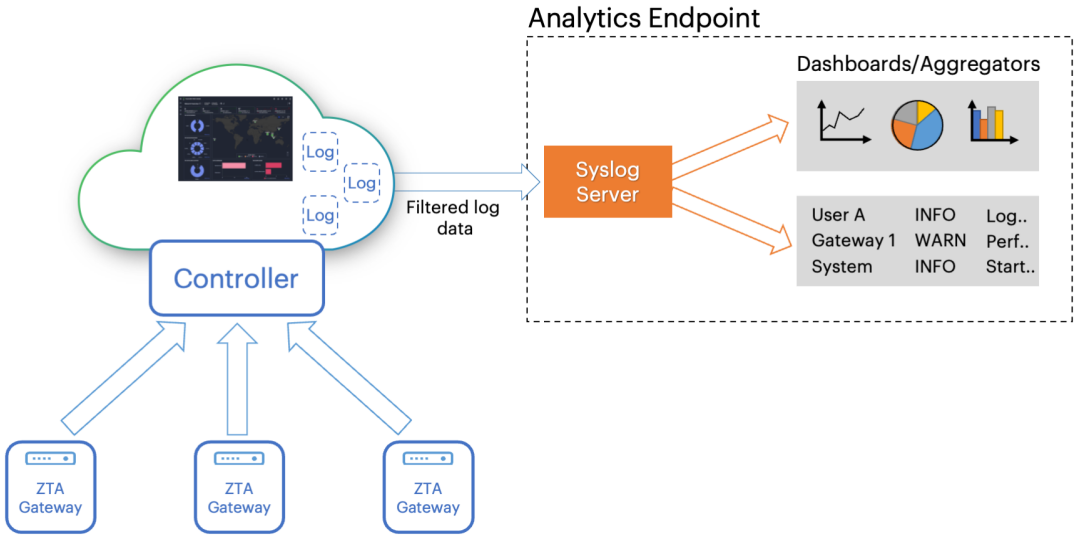

Exporting log data to an external analytics endpoint

To receive the log data exported from *nZTA*, your analytics endpoint should employ a *syslog* service that supports ingestion of data in JSON or WELF format. Syslog is a protocol defined through RFC 5424, and *nZTA* supports exporting log data to any Security Information and Event Management (SIEM) system that can accept and parse syslog messages. Supported configurations are:

- A publicly-reachable syslog server for your *Controller*. This requires you to configure *nZTA* with the hostname or IP address and port on which the analytics service is listening, and with the client certificate to use to authenticate *nZTA* to the syslog service.
- An *on-prem* syslog server for your *nZTA Gateway*(s). This requires you to configure *nZTA* with the hostname or IP address and port on which the analytics service is listening, and with the required protocol for the syslog service.

> *nSA* exports log data at 30 minutes intervals.

To use this facility, you must:

- Obtain and import a trusted Server Certificate Authority (CA) certificate suitable for the analytics service, see Importing a Trusted Server CA Certificate.
- Obtain and import a client authentication certificate for *nZTA*, see Adding a Client Certificate to the Controller.

> A client certificate is not required for an on-prem syslog server.

- Configure the *Controller* with the details of either:
  - Your publicly-reachable syslog server, see Adding a Public Syslog Server to the Controller.
  - Your on-prem Gateway syslog server, see Adding an On-Prem Syslog Server for nZTA Gateways.

Importing a Trusted Server CA Certificate

To export your logs to an analytics service, you first need to add a certificate for a Trusted Server Certificate Authority (Server CA).

Make sure you have a suitable certificate file, in PEM (Base64 ASCII encoded) format, stored on your local workstation before starting this procedure.

> *nSA* supports certificate files in PEM format only. Other formats, such as PFX, are not currently supported.

To import a certificate for a trusted server CA:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nSA* menu, click the **Administration** icon, then select **Enterprise Integration > Trusted Server CA**.

   The *Trusted Server CAs* page appears.

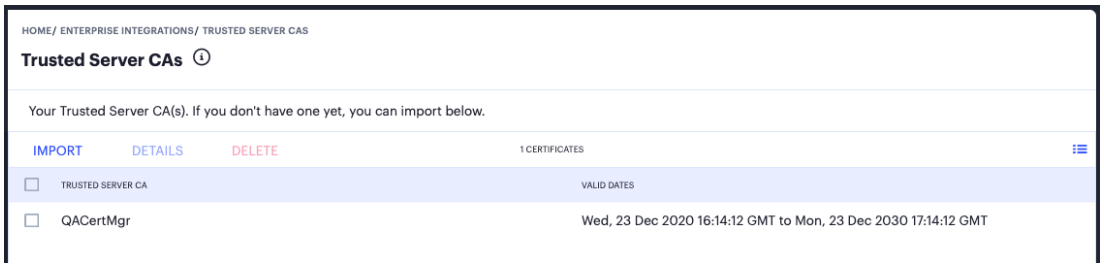

Viewing Trusted Server CA certificates

3. To import a Server CA certificate, click **Import**.

   The *Import Server CA* page appears.

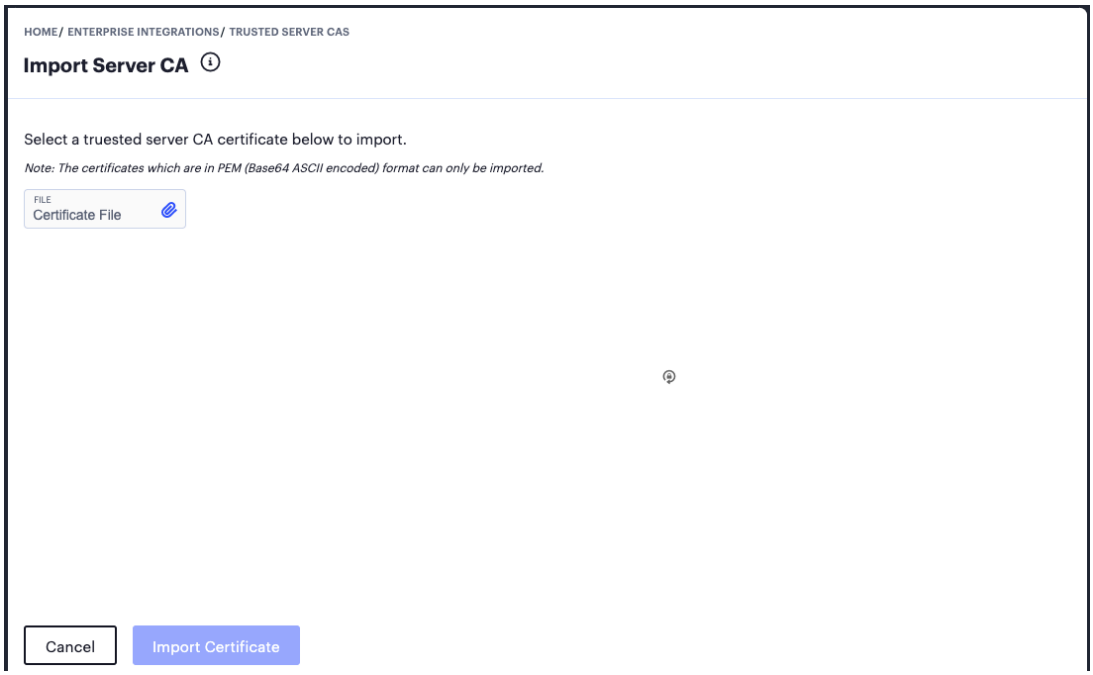

   Importing a Trusted Server CA certificate

4. Click **Certificate File** and use the subsequent dialog to locate the certificate file from your local workstation file system.

5. To import the selected certificate, click **Import Certificate**.

   The *Trusted Server CAs* page appears showing the successfully imported certificate.

6. (Optional) To view the details stored in a certificate, select the certificate checkbox and click **Details**.

7. (Optional) To delete a certificate, select the certificate checkbox and click **Delete**.

   You can import more than one trusted server CA certificate.

After you have imported your trusted server CA certificate to the *Controller*, proceed to add your client certificate (see Adding a Client Certificate to the Controller).

Adding a Client Certificate to the *Controller*

Before you can export your logs to an analytics service, you must add a client certificate to verify the identity of the *Controller* to the external service.

> A client certificate is not required for an on-prem syslog server.

Make sure you have a suitable client certificate file, in PEM (Base64 ASCII encoded) format, stored on your local workstation before starting this procedure.

> *nSA* supports certificate files in PEM format only. Other formats, such as PFX, are not currently supported.

To import a client certificate:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nSA* menu, click the **Administration** icon, then select **Enterprise Integration > Client Authentication**.

   The *Client Authentication* page appears.

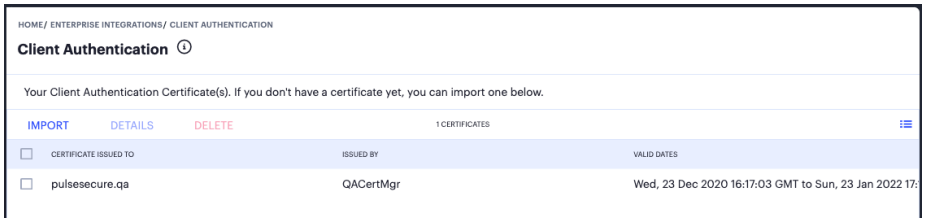

Viewing Client Authentication Certificates

3. To import a client certificate, click **Import**.

   The *Import Certificate* page appears.

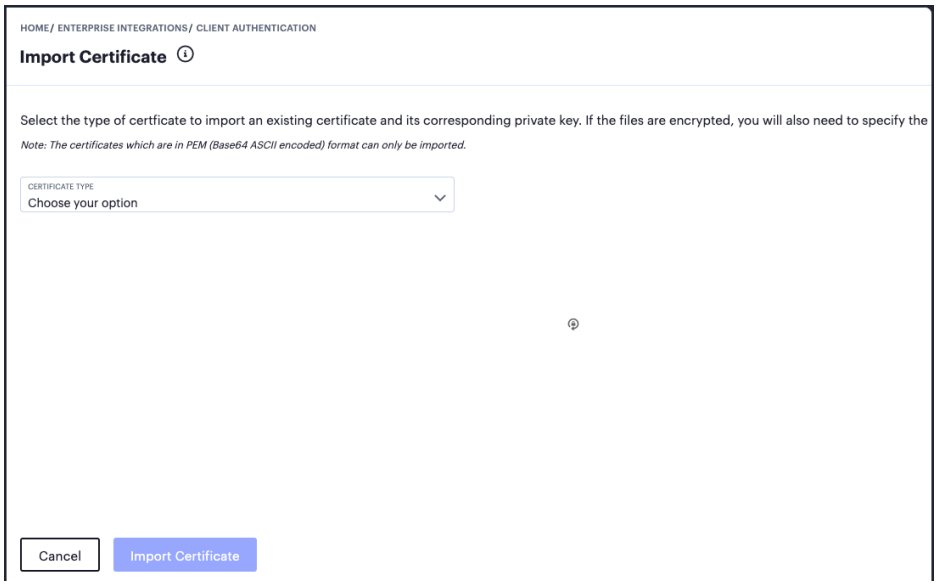

   Importing a Client certificate

4. Click **Certificate Type** and select one of the following options:

   - **Client Certificate with embedded private key**: Use this option if you are importing a certificate file with an embedded private key.
   - **Client Certificate with separate private key**: Use this option if you are importing a certificate file with a separate private key file.

   nSA updates the page to show the relevant import fields based on your selected option.

5. Click **Certificate File** and use the subsequent dialog to locate the certificate file from your local workstation file system.

6. (Optional) If you selected to use a separate private key, click **Private Key File** and use the subsequent dialog to locate the private key file from your local workstation file system.

7. (Optional) For **Password Key**, enter your private key passphrase.

8. To import the selected client certificate, click **Import Certificate**.

   The *Client Authentication* page appears showing the successfully imported client certificate.

9. (Optional) To view the details stored in a certificate, select the certificate checkbox and click **Details**.
10. (Optional) To delete a certificate, select the certificate checkbox and click **Delete**.

> You can import more than one client certificate.

After you have imported your Trusted Server CA certificate and Client certificate to the *Controller*, proceed to add the details of your external analytics service (see Adding a Public Syslog Server to the Controller).

Adding a Public Syslog Server to the *Controller*

Before you configure a publicly-reachable syslog server on the *Controller*, ensure that:

- Your syslog server is publicly-reachable, and that you know the hostname or IP address and port of the service.
- You have already added suitable Server CA and Client certificates.

To configure a syslog server in the *Controller*:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nSA* menu, click the **Administration** icon, then select **Enterprise Integration > Syslog Servers**.

   The *Enterprise Integrations* page appears.

   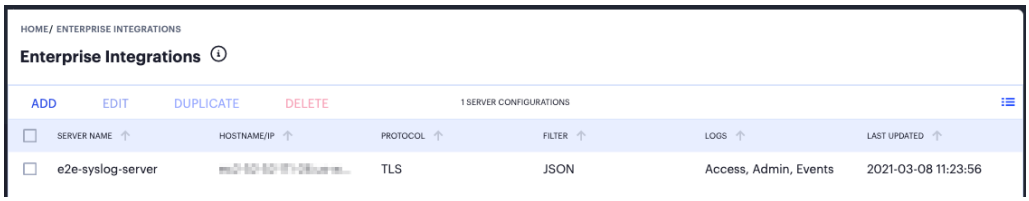

   Viewing Enterprise Integration Syslog Servers

3. Click **Add**.

   The *Add New Configuration* page appears, on the *Configuration* tab.

4. For **Type**, select *Syslog via Controller*.

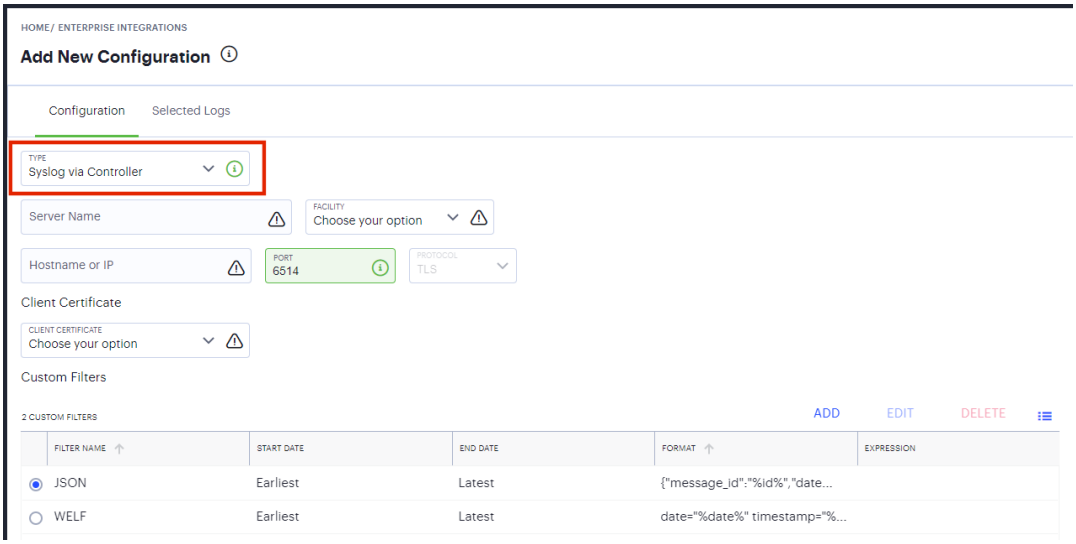

Add a new Syslog Server configuration

5. Enter data for the following fields:

   - **Server Name**: An identifying name for this syslog server configuration.
   - **Facility**: The syslog facility level *nSA* should use while exporting log data.
   - **Hostname or IP**: The hostname or IP address of the syslog server. This must match the value contained in the *Subject Alternative Name* specified in the Server CA certificate applicable to this configuration.
   - **Port**: The port on which the syslog server is listening.
   - **Protocol**: *This field is read-only.nZTA* supports only the TLS protocol in this instance.
   - **Client Certificate**: The client certificate you want to use with this syslog server.

6. Specify or create a **Custom Filter** to select the log data sent to the syslog server. Choose from:

   - Use either of the predefined "JSON" or "WELF" filters. These do not place restrictions on the log data and forward all logs to the external syslog server, using the respective data format.
   - Create your own filter, based on log parameter rules you define yourself, see Setting a Custom Syslog Filter.

   The predefined JSON and WELF filters are fixed and cannot be updated or deleted.

7. Click **Next** to select the required log files.

   The *Selected Logs* tab appears.

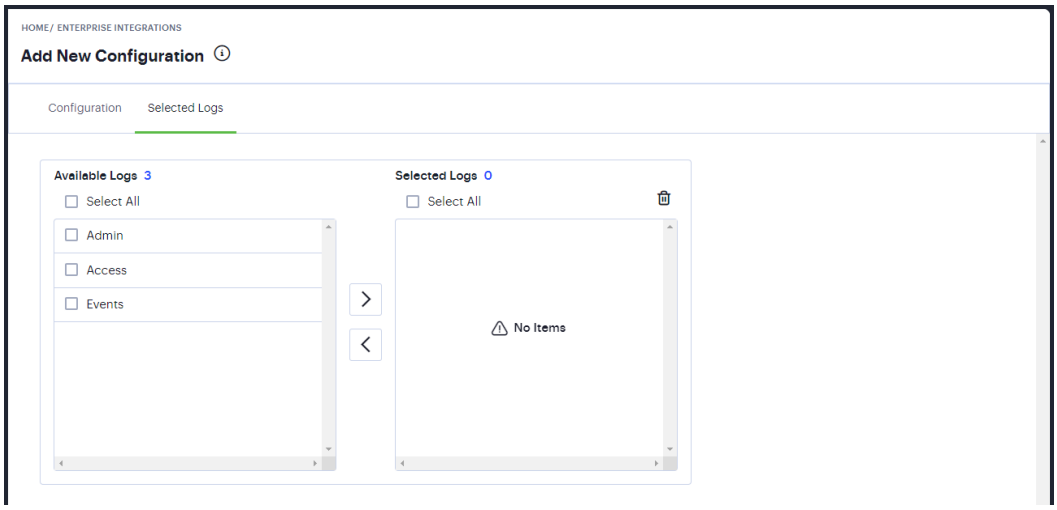

   Select the logs to be exported

8. From the list of *Available Logs*, select the check box adjacent to the desired log and click the arrow to move it to the the *Selected Logs* list.

   Repeat this step for each log you want to include. Use the reverse process to remove included logs.

9. To create a syslog server configuration with the current settings, click **Save Changes**.

Optionally:

- To edit an existing syslog server configuration, select the check box adjacent to the configuration entry and click **Edit**.

- To duplicate an existing syslog server configuration, select the check box adjacent to the configuration entry and click **Duplicate**. This creates a complete copy of the chosen configuration using the same server name appended with "-copy".
- To delete an existing syslog server configuration, select the check box adjacent to the configuration entry and click **Delete**.

Adding an On-Prem Syslog Server for *nZTA Gateways*

Before you configure an on-prem syslog server for *nZTA Gateways* to *Controller*, make sure you have added a suitable Server CA certificate. To learn more, see Importing a Trusted Server CA Certificate.

Ensure that know the hostname or IP address and port of the service.

To configure an on-prem syslog server for *nZTA Gateways* in the *Controller*:

1. Log into the *Controller* as a Tenant Admin, see Logging in as a Tenant Administrator.
2. From the *nSA* menu, click the **Administration** icon, then select **Enterprise Integration > Syslog Servers**.

   The *Enterprise Integrations* page appears.

   Viewing Enterprise Integration Syslog Servers

3. Click **Add**.

   The *Add New Configuration* page appears, on the *Configuration* tab.

4. For **Type**, select *Syslog from Gateways*.

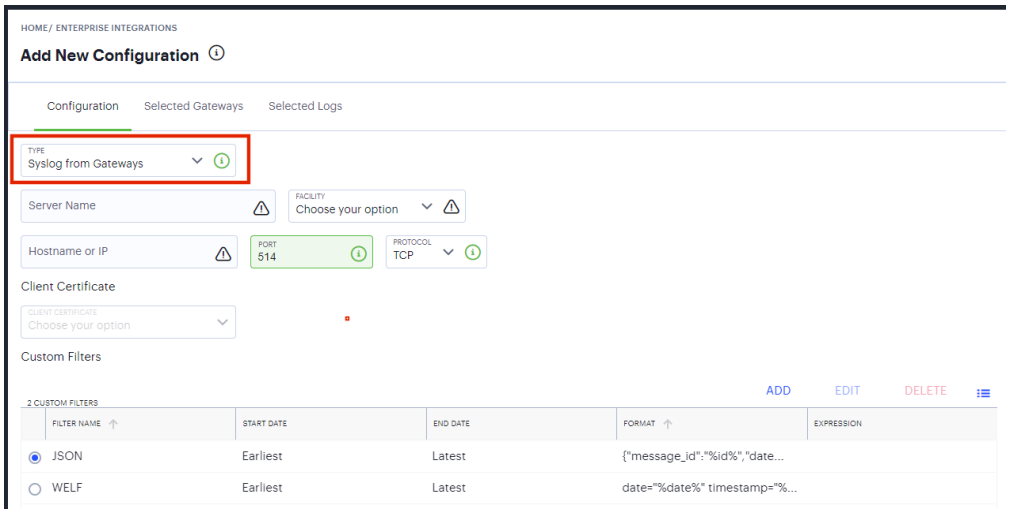

Add a new Syslog Server configuration

5. Enter data for the following fields:

   - **Server Name**: An identifying name for this syslog server configuration.
   - **Facility**: The syslog facility level *nSA* should use while exporting log data.
   - **Hostname or IP**: The hostname or IP address of the syslog server. This must match the value contained in the *Subject Alternative Name* specified in the Server CA certificate applicable to this configuration.
   - **Port**: The port on which the syslog server is listening.
   - **Protocol**: the required protocol for communicating with your on-prem syslog server. Both *TCP* and *UDP* are supported.
   - **Client Certificate**: This field is unavailable.

6. Specify or create a **Custom Filter** to select the log data sent to the syslog server. Choose from:

    - Use either of the predefined "JSON" or "WELF" filters. These do not place restrictions on the log data and forward all logs to the external syslog server, using the respective data format.
    - Create your own filter, based on log parameter rules you define yourself, see Setting a Custom Syslog Filter.

    The predefined JSON and WELF filters are fixed and cannot be updated or deleted.

7. Click **Next** to select the required gateways.

    The *Selected Gateways* tab appears.

    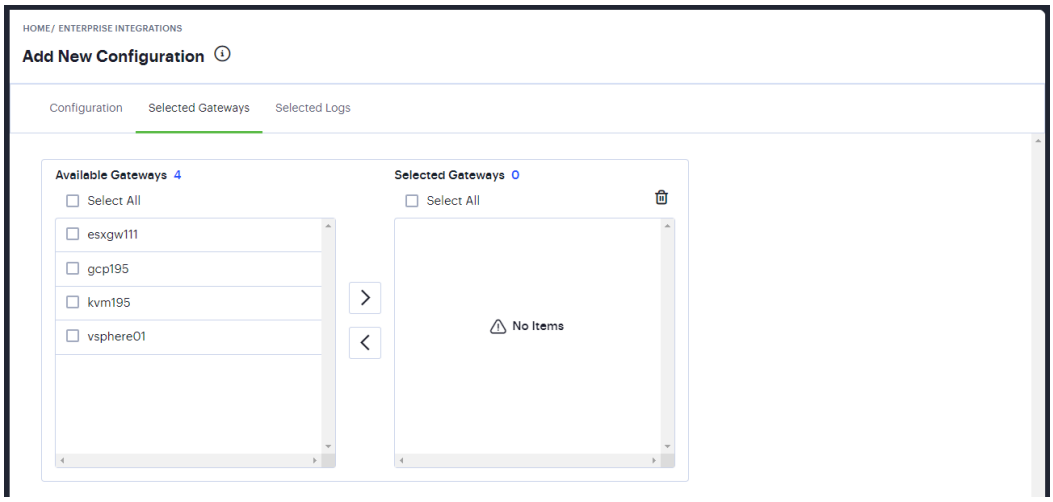

    Select the gateways

8. From the list of *Available Gateways*, select the checkbox adjacent to the desired Gateway and click the arrow to move it to the the *Selected Logs* list.

    Repeat this step for each Gateway you want to include. Use the reverse process to remove included *nZTA Gateways*.

9. Click **Next** to select the required log files.

   The *Selected Logs* tab appears.

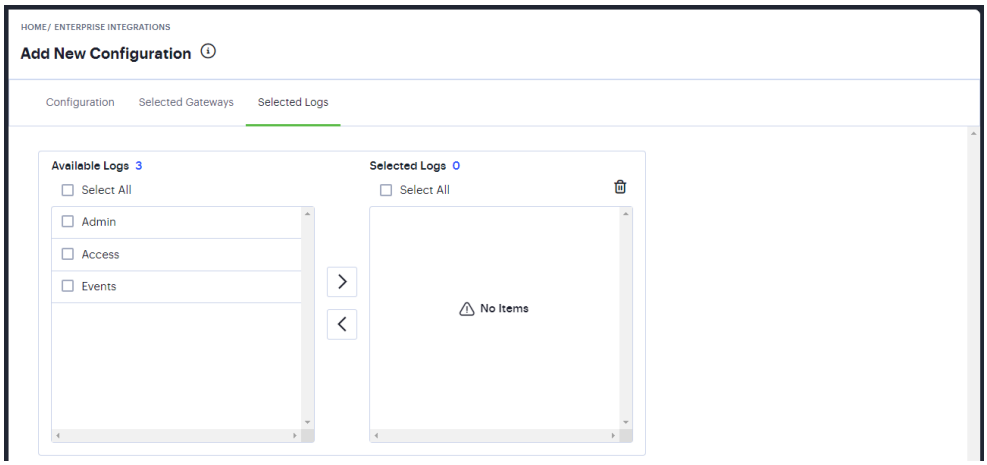

   Select the logs to be exported

10. From the list of *Available Logs*, select the checkbox adjacent to the desired log and click the arrow to move it to the the *Selected Logs* list.

    Repeat this step for each log you want to include. Use the reverse process to remove included logs.

11. To create a syslog server configuration with the current settings, click **Save Changes**.

Optionally:

- To edit an existing syslog server configuration, select the checkbox adjacent to the configuration entry and click **Edit**.
- To duplicate an existing syslog server configuration, select the checkbox adjacent to the configuration entry and click **Duplicate**. This creates a complete copy of the chosen configuration using the same server name appended with "-copy".
- To delete an existing syslog server configuration, select the checkbox adjacent to the configuration entry and click **Delete**.

Setting a Custom Syslog Filter

nSA enables you to create a custom filter that builds a query to specify the data items exported to your syslog server. You can specify criteria based on matching log data fields, date selection, and log output format.

This procedure assumes you are in the process of creating or editing a syslog server configuration. To learn more, see Adding a Public Syslog Server to the Controller.

To add a custom syslog filter:

1. From the *Add New Configuration* page, locate the *Custom Filters* section and click **Add**.

   The *Custom Filter* dialog appears.

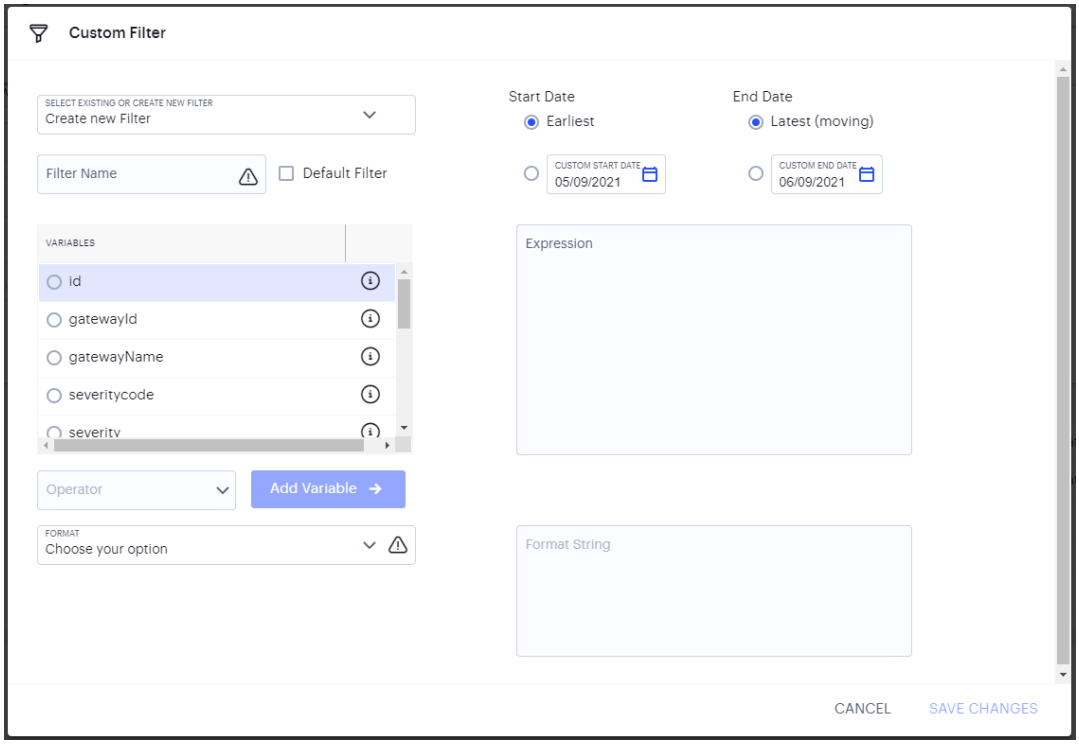

   Setting a custom syslog filter

2. Use **Select existing or create new filter** to either:

   - Create a new filter from scratch.
   - Select an existing filter as a template. This option pre-populates the **Expression** and **Format String** fields with values used in the selected filter, which you can then use to create your new filter.

3. Enter a **Filter Name**.

4. Select a **Start Date** and **End Date** for the log data to be included. Use the default *Earliest* and *Latest* values to include all data, or select custom start and end dates, or a combination.

   nZTA stores log data for the previous 30 days.

5. (Optional) To use this filter as the default for your Syslog Server configurations, select **Default Filter**.

6. (Optional) Select from the list of available log data **Variables**, select an **Operator**, and then click **Add Variable** to add a filter based on the chosen value.

   nSA populates the **Expression** field with an expression matching your selection. An example *value* is added to the right of the operator, but this can be freely edited to your required value.

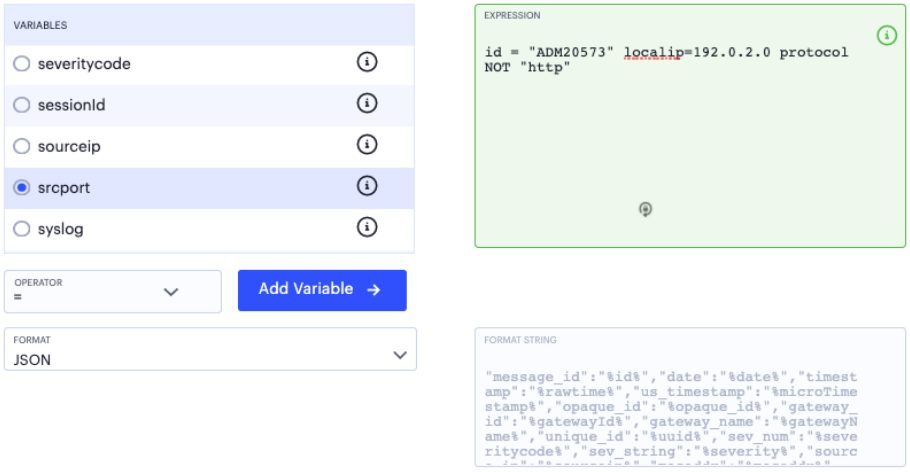

   Building a custom syslog filter expression based on the log fields: message ID, local IP address, and protocol.

   Hover your pointer over the (i) icon to the right of each variable to view a tooltip showing example usage.

7. (Optional) Repeat the previous step as required for each log data variable you want to include in the syslog filter.

8. Select an output **Format** for the log line. Choose from JSON or WELF. The resultant formatted string is shown in the box provided.

9. To save your filter settings, click **Save Changes**.\

   The Add New Configuration screen appears, showing your new filter as selected. Continue to create your syslog server configuration.

APPENDIX: Applications Supported by *nSA*

Cloud/SaaS/On-Premises Applications

Applications
Office 365 with On-Premises AD
Salesforce
Box
G-Suite
Zoom
SAP-Analytics Cloud
Zendesk
WebEx
Slack
JIRA
Confluence
BambooHR
GitHub
PagerDuty
Yammer
Okta