



Pulse Secure Services Director Getting Started Guide

Supporting Pulse Secure Services Director 19.1

Product Release	19.1
Published	20 May 2019
Document Version	1.0

Pulse Secure, LLC
2700 Zanker Road,
Suite 200 San Jose
CA 95134

<https://www.pulsesecure.net>

© 2019 Pulse Secure, LLC. All rights reserved. Pulse Secure and the Pulse Secure logo are trademarks of Pulse Secure, LLC in the United States.

All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Pulse Secure, LLC. assumes no responsibility for any inaccuracies in this document. Pulse Secure, LLC. reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Pulse Secure Services Director Getting Started Guide

The information in this document is current as of the date on the title page.

END USER LICENSE AGREEMENT

The Pulse Secure product that is the subject of this technical documentation consists of (or is intended for use with) Pulse Secure software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.pulsesecure.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Contents

PREFACE	1
DOCUMENT CONVENTIONS	1
TEXT FORMATTING CONVENTIONS	1
COMMAND SYNTAX CONVENTIONS.....	2
REQUESTING TECHNICAL SUPPORT	2
SELF-HELP ONLINE TOOLS AND RESOURCES	2
OPENING A CASE WITH PSGSC	3
ABOUT THIS DOCUMENT	5
SERVICES DIRECTOR VA OVERVIEW.....	5
USING THE GETTING STARTED GUIDE.....	6
PREPARING TO INSTALL THE SERVICES DIRECTOR VIRTUAL APPLIANCE.....	9
OVERVIEW: PLATFORMS	9
PREREQUISITES.....	9
REQUIRED SOFTWARE FOR INSTALLATION	10
REQUIRED HARDWARE RESOURCES FOR VIRTUAL MACHINES	10
REQUIRED FILES AND INFORMATION.....	11
CRITICAL PORTS THAT MUST BE OPEN.....	12
INSTALLING THE SERVICES DIRECTOR VA ON VSPHERE.....	13
OVERVIEW: SERVICES DIRECTOR VA ON VSPHERE.....	13
OBTAINING THE SERVICES DIRECTOR VA OVA PACKAGE	14
OBTAINING SERVICES DIRECTOR LICENSES.....	14
CREATING A VM IN VSPHERE	14
ACCESSING THE SERVICES DIRECTOR VA ON VMWARE.....	15
INSTALLING THE SERVICES DIRECTOR VA ON KVM-QEMU	17
OVERVIEW: SERVICES DIRECTOR VA ON KVM-QEMU.....	17
OBTAINING THE SERVICES DIRECTOR VA KVM IMAGE	18
OBTAINING SERVICES DIRECTOR LICENSES.....	18
CREATING THE SERVICES DIRECTOR VA ON A KVM SERVER	18
CREATING A VM USING THE LIBVIRT COMMAND LINE INTERFACE.....	19
CREATING A VM USING THE VMM GRAPHICAL USER INTERFACE	20
ACCESSING THE SERVICES DIRECTOR VA ON KVM	25

INSTALLING THE SERVICES DIRECTOR VA ON AMAZON WEB SERVICES	27
OVERVIEW: SERVICES DIRECTOR VA ON AMAZON WEB SERVICES.....	27
OBTAINING SERVICES DIRECTOR LICENSES.....	27
LAUNCHING AND CONFIGURING THE PRIMARY SERVICES DIRECTOR ON AWS.....	28
PREPARING AWS INFRASTRUCTURE	28
PREPARING AN AWS SECURITY GROUP	40
LAUNCHING A SERVICES DIRECTOR AMI INSTANCE ON AWS.....	64
CREATING ELASTIC IP ADDRESSES FOR THE SERVICES DIRECTOR INSTANCE	69
UPDATING SECURITY RULES FOR SERVICES DIRECTOR INSTANCE IP ADDRESSES	74
RETRIEVING THE DEFAULT PASSWORD FOR A SERVICES DIRECTOR INSTANCE	74
ACCESSING YOUR SERVICES DIRECTOR INSTANCE FOR THE FIRST TIME	77
LAUNCHING AND CONFIGURING THE SECONDARY SERVICES DIRECTOR ON AWS.....	78
RUNNING THE SERVICES DIRECTOR VA SETUP WIZARD.....	79
OVERVIEW: SETUP WIZARD.....	79
INSTALLING AND CONFIGURING A PRIMARY SERVICES DIRECTOR.....	80
STARTING THE SETUP WIZARD.....	80
DEFINING A SERVICE ENDPOINT ADDRESS.....	89
REDEEMING A LICENSE TOKEN	92
GENERATING A SELF-SIGNED SSL CERTIFICATE.....	93
ADDING CERTIFICATES AND LICENSES.....	95
COMPLETING THE SERVICES DIRECTOR INSTALLATION	103
INSTALLING AND CONFIGURING A SECONDARY SERVICES DIRECTOR.....	105
ACCESSING YOUR SERVICES DIRECTOR VA.....	110
UPDATING SERVICES DIRECTOR VA SETTINGS.....	113
OVERVIEW: SERVICES DIRECTOR VA SETTINGS.....	113
UPDATING GENERAL SETTINGS	113
UPDATING MONITORING SETTINGS	115
UPDATING METERING SETTINGS.....	115
UPDATING LICENSING SETTINGS.....	116
UPDATING LOGGING SETTINGS.....	116
UPDATING DEPLOYMENT SETTINGS	117
UPDATING BANDWIDTH LICENSING SETTINGS.....	117
UPDATING CONTROLLER LICENSING SETTINGS	117
UPDATING INSTANCE REGISTRATION SETTINGS.....	117
UPDATING TELEMETRY SETTINGS	117
UPDATING METERING ALERTS AND NOTIFICATIONS SETTINGS	118
CONFIGURING THE FLA CHECKER.....	118
UPDATING AUTO CLEANUP OF FAILED VTMS	118
UPDATING DATE AND TIME SETTINGS.....	118

UPDATING ADMINISTRATION CREDENTIALS.....	119
UPDATING EMAIL SETTINGS	119
UPDATING THE SSL CERTIFICATE.....	119
UPDATING THE REST API PORT	120
UPDATING SECURITY SETTINGS	120
CHANGING THE MASTER PASSWORD FOR THE SERVICES DIRECTOR VA.....	121
CHANGING THE MASTER PASSWORD	121
ADDING VIRTUAL TRAFFIC MANAGERS TO THE SERVICES DIRECTOR.....	125
OVERVIEW: ADDING VIRTUAL TRAFFIC MANAGERS TO THE SERVICES DIRECTOR	125
WORKING WITH VTM COMMUNICATIONS CHANNEL	126
ENABLING A VTM CLUSTER TO OPERATE BEHIND A NAT DEVICE	126
DISABLING COMMS CHANNEL ON A VTM	127
ADDING RESOURCES REQUIRED FOR VIRTUAL TRAFFIC MANAGERS.....	128
ADDING A LICENSE TO THE SERVICES DIRECTOR.....	128
ADDING A FEATURE PACK TO THE SERVICES DIRECTOR	129
ADDING AN OWNER TO THE SERVICES DIRECTOR.....	141
ADDING A LEGACY FLA LICENSE TO THE SERVICES DIRECTOR	144
ADDING AN AUTO-ACCEPT POLICY TO THE SERVICES DIRECTOR.....	147
ADDING A CLOUD REGISTRATION RESOURCE TO THE SERVICES DIRECTOR	149
REGISTERING AN EXTERNALLY-DEPLOYED VIRTUAL TRAFFIC MANAGER.....	153
PREPARING TO REGISTER A VIRTUAL TRAFFIC MANAGER (UNIVERSAL FLA).....	153
REGISTERING A VIRTUAL TRAFFIC MANAGER (UNIVERSAL FLA).....	154
PREPARING TO REGISTER A VIRTUAL TRAFFIC MANAGER (LEGACY FLA LICENSE)	160
REGISTERING A VIRTUAL TRAFFIC MANAGER (INACTIVE REST API).....	161
REGISTERING A VIRTUAL TRAFFIC MANAGER (PRE-10.1 VTM SOFTWARE VERSION)	164
SELF-REGISTERING AN EXTERNALLY-DEPLOYED VIRTUAL TRAFFIC MANAGER.....	169
OVERVIEW: VTM SELF-REGISTRATION (VMWARE).....	169
REQUESTING SELF-REGISTRATION DURING VTM INSTALLATION	172
REQUESTING SELF REGISTRATION ON A CONFIGURED VTM	177
VIEWING VTM INSTANCE REGISTRATION REQUESTS	178
PROCESSING SELF-REGISTRATION REQUESTS MANUALLY	182
REQUESTING RE-REGISTRATION OF A VTM	185
SELF-REGISTERING A CLOUD-BASED VIRTUAL TRAFFIC MANAGER.....	186
OVERVIEW: VTM SELF-REGISTRATION (CLOUD)	186
CREATING A CLOUD-BASED VIRTUAL TRAFFIC MANAGER.....	188
WORKING WITH VIRTUAL TRAFFIC MANAGERS	193
OVERVIEW: WORKING WITH VIRTUAL TRAFFIC MANAGERS	193
VIEWING VIRTUAL TRAFFIC MANAGERS	194
UNDERSTANDING BASIC DETAILS OF A VIRTUAL TRAFFIC MANAGER.....	194

UNDERSTANDING LIFECYCLE STATUS (EXTERNALLY-DEPLOYED VTMS)	195
UNDERSTANDING LIFECYCLE STATUS (DEPLOYED VTMS).....	196
UNDERSTANDING THE INSTANCE HEALTH OF A VIRTUAL TRAFFIC MANAGER.....	196
UNDERSTANDING THE LICENSING HEALTH OF A VIRTUAL TRAFFIC MANAGER.....	197
VIEWING FULL DETAILS FOR A VIRTUAL TRAFFIC MANAGER	198
CHANGING THE DISPLAY ORDER OF VTMS.....	199
FILTERING VTMS	200
UPDATING DETAILS FOR A VIRTUAL TRAFFIC MANAGER	201
UNDERSTANDING VSERVER STATUS.....	201
DELETING A VIRTUAL TRAFFIC MANAGER	203
CONFIGURING AUTO CLEANUP OF VIRTUAL TRAFFIC MANAGERS	204
EXAMPLE OF AUTO CLEANUP	205
WORKING WITH APPLICATION TEMPLATES (ENTERPRISE FEATURE TIER)	207
OVERVIEW OF APPLICATION TEMPLATES AND TEMPLATE INSTANCES	207
ADDING AN APPLICATION TEMPLATE TO SERVICES DIRECTOR.....	210
CREATING AND APPLYING A TEMPLATE INSTANCE.....	212
EDITING A TEMPLATE INSTANCE	217
REMOVING A VTM APPLICATION BY DELETING A TEMPLATE INSTANCE	218
RELICENSING VIRTUAL TRAFFIC MANAGERS	221
PREPARING TO RELICENSE A VIRTUAL TRAFFIC MANAGER (LEGACY FLA TO UNIVERSAL FLA)	221
RELICENSING A VIRTUAL TRAFFIC MANAGER INSTANCE	221
PROCESSING VIRTUAL TRAFFIC MANAGER METERING DISCREPANCY WARNINGS.....	224
UNDERSTANDING METERING DISCREPANCY WARNINGS.....	225
WORKING WITH VIRTUAL TRAFFIC MANAGER CLUSTERS.....	229
OVERVIEW: WORKING WITH VIRTUAL TRAFFIC MANAGER CLUSTERS.....	229
UNDERSTANDING VIRTUAL TRAFFIC MANAGER CLUSTER DETAILS.....	230
CREATING A VIRTUAL TRAFFIC MANAGER CLUSTER	233
UPDATING A VIRTUAL TRAFFIC MANAGER CLUSTER.....	234
WORKING WITH VTM CLUSTER BACKUPS.....	237
OVERVIEW: VTM CLUSTER BACKUPS	237
CREATING A CLUSTER BACKUP SCHEDULE.....	239
UPDATING A CLUSTER BACKUP SCHEDULE	241
ADDING A BACKUP SCHEDULE TO A CLUSTER.....	242
VIEWING BACKUPS FOR A CLUSTER.....	243
UPDATING DETAILS FOR A CLUSTER BACKUP	245
PERFORMING AN IMMEDIATE BACKUP FOR A CLUSTER.....	246
COMPARING TWO CLUSTER BACKUPS	248
RESTORING A BACKUP TO A CLUSTER.....	250
UPLOADING A CLUSTER BACKUP TO A VIRTUAL TRAFFIC MANAGER.....	253

DELETING A CLUSTER BACKUP	256
MOVING A VTM BETWEEN CLUSTERS.....	257
DELETING AN EMPTY VIRTUAL TRAFFIC MANAGER CLUSTER	258
WORKING WITH USER AUTHENTICATION	259
OVERVIEW: VTM USER AUTHENTICATION.....	259
OVERVIEW: SERVICES DIRECTOR USER AUTHENTICATION	260
CREATING AN AUTHENTICATOR	261
VIEWING AUTHENTICATORS	261
CREATING AN LDAP AUTHENTICATOR	262
CREATING A RADIUS AUTHENTICATOR	264
CREATING A TACACS+ AUTHENTICATOR.....	266
CREATING A PERMISSION GROUP	268
VIEWING PERMISSION GROUPS.....	268
CREATING A PERMISSION GROUP (VTM USER AUTHENTICATION)	270
CREATING A PERMISSION GROUP (SD USER AUTHENTICATION).....	272
CREATING AN ACCESS PROFILE (VTM USER AUTHENTICATION ONLY)	272
VIEWING ACCESS PROFILES	273
CREATING AN ACCESS PROFILE	274
APPLYING USER AUTHENTICATION TO A VTM	275
WORKING WITH VTM TEMPLATES	277
WORKING WITH VTM ANALYTICS.....	279
OVERVIEW: VTM ANALYTICS (ENTERPRISE CUSTOMERS ONLY)	279
UNDERSTANDING THE ANALYTICS SYSTEM.....	280
CONFIGURING VTM ANALYTICS ON THE SERVICES DIRECTOR	281
UNDERSTANDING THE AUTOMATIC EXPORT OF VTM ANALYTICS DATA	283
QUERYING VTM ANALYTICS FROM THE SERVICES DIRECTOR.....	284
CREATING ANALYTICS RESOURCES.....	285
CREATING A LOG EXPORT TYPE.....	285
CREATING AN ANALYTICS PROFILE	287
ADDING ANALYTICS ENDPOINT RESOURCES TO THE SERVICES DIRECTOR	289
ENABLING ANALYTICS ON A VTM CLUSTER	297
ADDING AN ANALYTICS PROFILE TO A VTM CLUSTER.....	297
WORKING WITH ANALYTICS DATA ON THE SERVICES DIRECTOR.....	298
ACCESSING THE VADC ANALYTICS APPLICATION.....	299
RETURNING TO THE SERVICES DIRECTOR VA.....	301
CHOOSING A DATA METRIC.....	302
CHOOSING A TIME PERIOD	303
CHOOSING A SAMPLING RATIO	305
WORKING WITH THE COMPONENT FILTER	307

WORKING WITH THE EXTENDED FILTER	316
USING THE SANKEY DIAGRAM	325
USING THE TABLE GRAPH	332
USING CHARTS	334
USING THE DATASET VIEW.....	360
WORKING WITH THE LOGS VIEW.....	363
WORKING WITH HIGH AVAILABILITY.....	369
OVERVIEW: HIGH AVAILABILITY ON SERVICES DIRECTOR.....	369
CREATING A HIGH AVAILABILITY PAIR IN THE SERVICES DIRECTOR VA	371
VIEWING HIGH AVAILABILITY STATUS.....	372
TAKING A BACKUP OF YOUR SERVICES DIRECTOR	374
RESPONDING TO REPORTED HEALTH ISSUES	375
SWAPPING THE ROLES OF THE HA NODES.....	377
PERFORMING A FAILOVER FROM THE STANDBY NODE	377
EJECTING A NODE FROM AN HA PAIR.....	381
EJECTING A STANDBY NODE FROM THE ACTIVE NODE.....	382
RECOVERING FROM A FAILED ACTIVE NODE.....	385
TO PERFORM A FORCED FAILOVER FROM THE STANDBY NODE	386
RECOVERING FROM A SPLIT BRAIN SCENARIO	389
UNDERSTANDING HOW THE SPLIT BRAIN SCENARIO ARISES.....	389
VIEWING THE SPLIT BRAIN SCENARIO.....	390
RESOLVING A SPLIT BRAIN SCENARIO.....	391
CONVERTING AN EJECTED NODE INTO A STANDALONE ACTIVE NODE	395
CONVERTING AN UPGRADED NODE INTO A STANDALONE ACTIVE NODE.....	398
RECOVERING FROM A SERVICES DIRECTOR FAILURE	401
OVERVIEW: RECOVERING FROM A SERVICES DIRECTOR FAILURE	401
UNDERSTANDING A BACKUP FILE	403
CONFIGURING A SCHEDULED BACKUP SCHEDULE.....	403
CONFIGURING THE BACKUP SCHEDULE	404
UPDATING THE BACKUP SCHEDULE	407
RESTORING A SERVICES DIRECTOR FROM A LOCAL BACKUP.....	409
RESTORING A SERVICES DIRECTOR FROM A REMOTE BACKUP	410
RESTORING A SERVICES DIRECTOR USING THE SETUP WIZARD	412
STARTING AND STOPPING THE SERVICES DIRECTOR SERVICE.....	418
RESTARTING THE SERVICES DIRECTOR VA	418
ENTERING THE MASTER PASSWORD AFTER A VIRTUAL MACHINE RESTART	418
CREATING SERVICES DIRECTOR REPORTS	423
VIEWING REPORTS AND DIAGNOSTICS.....	423

THE VTM INSTANCE ALLOCATION REPORT424

THE BANDWIDTH ALLOCATION REPORT.....426

THE CPU UTILIZATION REPORT428

THE THROUGHPUT UTILIZATION REPORT.....429

VIEWING LOGS AND GENERATING SYSTEM DUMPS.....430

 VIEWING SYSTEM LOGS430

 GENERATING SYSTEM DUMPS.....431

WORKING WITH METERING LOGS432

 GENERATING METERING LOGS432

 DOWNLOADING METERING LOGS.....432

 DELETING METERING LOGS.....433

 MONITORING THE STORAGE CAPACITY OF METERING LOGS433

 CONFIGURING THE PHONE HOME FUNCTION.....435

 MANUALLY PHONING HOME A METERING LOG FILE435

 UNDERSTANDING METERING LOGS435

Preface

- **Document Conventions** 1
- **Requesting Technical Support** 2

Document Conventions

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

Text Formatting Conventions

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

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

Command Syntax Conventions

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

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

Requesting Technical Support

Technical product support is available through the Pulse Secure Global Support Center (PSGSC). If you have a support contract, then file a ticket with PSGSC.

- Product warranties—For product warranty information, visit <http://www.pulsesecure.net>.

Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.pulsesecure.net/support>
- Search for known bugs: <http://www.pulsesecure.net/support>
- Find product documentation: <https://www.pulsesecure.net/techpubs>
- Find solutions and answer questions using our Knowledge Base: <http://www.pulsesecure.net/support>
- Download the latest versions of software and review release notes: <http://www.pulsesecure.net/support>

- Search technical bulletins for relevant hardware and software notifications:
<http://www.pulsesecure.net/support>
- Open a case online in the CSC Case Management tool: <http://www.pulsesecure.net/support>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool:
<http://www.pulsesecure.net/support>

Opening a Case with PSGSC

You can open a case with PSGSC on the Web or by telephone.

- Use the Case Management tool in the PSGSC at <http://www.pulsesecure.net/support>.
- Call 1- 844-751-7629 (toll-free in the USA).

For international or direct-dial options in countries without toll-free numbers, see
<http://www.pulsesecure.net/support>.

About This Document

- **Services Director VA Overview 5**
- **Using the Getting Started Guide..... 6**

Services Director VA Overview

The Services Director Virtual Appliance (Services Director VA) enables you to configure and manage the Services Director as a virtual appliance. The Services Director VA provides a graphical user interface (GUI) that enables you to:

- License your Virtual Traffic Manager (vTM) instances.
- Register externally-deployed vTM instances.
- Configure the use of an external instance host by the Services Director.
- Deploy vTM instances using a configured instance host.
- Deploy cloud-based vTM instances on AWS.
- Transition deployed vTM instances through a lifecycle.
- Start, stop and restart your Services Director service.
- Protect your instance configurations (on a cluster basis) by taking automated and manual backups.
- Protect your Services Director configuration using a backup system.
- Protect your vTM passwords using encryption based on a Master Password.
- Perform health and monitoring reporting.
- Configure vTM analytics for a vTM cluster, and view resulting analytics graphs.
- Perform usage metering.
- Generate system logs and system dumps.

Note: Support for individual functions depends on your license type.

Note: The GUI is the main interface for the Services Director VA. However, a Command-Line Interface (CLI) is also included. The CLI is described in the *Secure Services Director Command Reference*.

Using the Getting Started Guide

This guide takes you through the installation, configuration and use of your Services Director VA.

The structure of this guide is as follows:

- **“Preparing to Install the Services Director Virtual Appliance” on page 9:** Describes the general Services Director VA installation process. It references platform-specific processes across a number of chapters:
 - **“Installing the Services Director VA on vSphere” on page 13.**
 - **“Installing the Services Director VA on KVM-QEMU” on page 17.**
 - **“Installing the Services Director VA on Amazon Web Services” on page 27.**
 - **“Running the Services Director VA Setup Wizard” on page 79.**
 - **“Updating Services Director VA Settings” on page 113.**
- **“Adding Virtual Traffic Managers to the Services Director” on page 125:** Describes the process of adding externally-deployed vTM instances to the Services Director VA. This includes manual registrations, the processing of self-registration requests, and the creation of cloud-based vTM instances.

Note: The installation and configuration of an instance host, and the deployment of vTM instances is described in the *Pulse Services Director Advanced User Guide*.

- **“Working with Virtual Traffic Managers” on page 193:** Describes how vTM instances are represented in the Services Director VA, methods for affecting this representation, and the lifecycle of externally-deployed vTM instances.

Note: The operation of traffic management and load balancing on individual vTM instances is not addressed by the Services Director. This requires use of a Pulse Secure Virtual Traffic Manager for each vTM instance.

- **“Working with Virtual Traffic Manager Clusters” on page 229:** Describes how vTM clusters and backups are used by both vTMs and the Services Director VA.
- **“Working with User Authentication” on page 259:** Describes how to configure user authentication for both vTMs and the Services Director VA.
- **“Working with vTM Analytics” on page 279:** Describes how to configure vTM analytics on the Services Director VA, and how to use the various analytics graph types.
- **“Working with High Availability” on page 369:** Describes how to operate a High Availability (HA) pair of Services Director VA nodes. This includes monitoring of status, error conditions, and methods for returning your HA pair to operation.

- **“Recovering from a Services Director Failure” on page 401:** Describes how to preserve the configuration of an HA pair, and how to recover a saved configuration for an existing Services Director VA. This also includes how to create a new Services Director VA from a saved configuration.
- **“Creating Services Director Reports” on page 423:** Describes how to generate and extract output from your Services Director VA. This includes metering logs and system logs.

Preparing to Install the Services Director Virtual Appliance

- [Overview: Platforms](#) 9
- [Prerequisites](#) 9
- [Critical Ports That Must Be Open](#) 12

Overview: Platforms

The Services Director Virtual Appliance (VA) can be installed on a number of platforms. Each platform has prerequisites that must be met before you begin installation, see [“Prerequisites” on page 9](#).

You can install the Services Director VA as a Virtual Machine on the following platforms:

- VMware, see [“Installing the Services Director VA on vSphere” on page 13](#).
- KVM-QEMU, see [“Installing the Services Director VA on KVM-QEMU” on page 17](#).
- Amazon Web Services (AWS), see [“Installing the Services Director VA on Amazon Web Services” on page 27](#).

After the Services Director VA is installed as a VM/instance, you must:

- Run the Services Director Setup Wizard to configure the Services Director VA for use, see [“Running the Services Director VA Setup Wizard” on page 79](#).
- Review and update all Services Director settings, see [“Updating Services Director VA Settings” on page 113](#).

Prerequisites

Before you install the Services Director VA and run the Setup Wizard, you must make sure that you have the correct software, files and configuration information.

Required Software for Installation

You need the following software to install the Services Director VA using a VMware hypervisor.

Software	Description
VMware vSphere ESXi 6.0+	Pulse Secure assumes that you are familiar with creating and managing VMs using vSphere. For detailed information about creating virtual machines using vSphere, refer to http://www.vmware.com/products/ .
Services Director VMware image in OVA format	This image is used to install the Services Director VA. You can obtain the Services Director OVA package from Pulse Secure Support.

You need the following software to install the Services Director VA using a KVM-QEMU hypervisor.

Software	Description
A virtualization toolset, such as libvirt or Virtual Machine Manager (VMM)	Pulse Secure assumes that you are familiar with creating and managing VMs using your chosen toolset. For detailed information about creating virtual machines on KVM-QEMU, refer to http://wiki.qemu.org/KVM .
Services Director KVM image in QCOW2 format	This image is used to install the Services Director VA on a KVM-QEMU hypervisor. You can obtain the Services Director KVM image in QCOW2 format from Pulse Secure Support.

You need an Amazon Web Services (AWS) account and a browser to use Services Director on AWS.

Required Hardware Resources for Virtual Machines

You need the following hardware resources to use Services Director VA on vSphere and KVM-QEMU.

VA Type	CPU	Memory	Disk
Services Director VA	4 vCPU	8 GB	46 GB

Your hardware must support the required configuration.

There are no hardware requirements for AWS, as it is cloud-based.

Required Files and Information

The following table lists the files and information required by the Services Director VA.

Note: All required files must be in accessible locations in your infrastructure during the installation process. For example, locate the files on an accessible server, or your local machine.

Information	Description
Hostnames	The hostname for the Services Director. When you are creating a High Availability pair, you will need a hostname for both the Primary and the Secondary Services Director nodes.
DNS Server	(Optional) The IP address for the primary name server. This is not required if you choose to configure your system using IP addresses rather than DNS hostnames. You can also specify a secondary name server if required.
Primary Address	The IP address for the Primary Services Director in a High Availability pair.
Secondary Address	The IP address for the Secondary Services Director in a High Availability pair.
Service Endpoint Address	The Management IP address for your High Availability Services Director installation. This IP address binds to the currently active Services Director.
SSL certification and private key	A self-signed Secure Socket Layer (SSL) certificate and private key file, which are used to protect and authenticate the REST API port. This is a local file or URL using HTTP, FTP, or SCP. For example: <code>scp://username:password@host/path/filename</code> Pulse Secure recommends that you do not use a CA-signed certificate.
Services Director License	The Services Director License, either for Cloud Service Providers or Enterprise customers. Note: If you have not received your Services Director License, contact Pulse Secure Support for assistance.
Resource Licenses	For Enterprise Services Director Licenses/Customers only. This includes Bandwidth Resource Licenses, and Analytics Resource Pack Licenses. Note: If you have not received your Licenses, contact Pulse Secure Support for assistance.
Add-On Licenses	An Add-On License is a historical license type, that is only supported on “old style” Services Director licenses. It is not compatible with “new style” Services Director licenses.
Legacy FLA License	(Optional) The Flexible Licensing Architecture (FLA) Legacy License is for: <ul style="list-style-type: none"> Any Virtual Traffic Manager (vTM) instances at version 10.0 or earlier. Any vTM instances that do not have an enabled REST API. vTMs that are at version 10.1 (or later) with their REST API enabled will use a pre-installed Universal License.

Information	Description
Administrator user and password	The administrator password for the Services Director. This password is used to access the Services Director GUI and CLI. The default administrator user is admin and the password is password.
SMTP server and port	(Optional) The hostname (or IP address if DNS is not configured) of the SMTP server and port. External DNS and external access for SMTP traffic is required for email notification of events and failures to function.
Email notification address	(Optional) A valid email address to which notification of events and failures are to be sent.

Critical Ports That Must Be Open

The following table lists ports must be open on the Services Director VA.

Port	Open to Connections From	Description	Protocol
22	Any machine that may legitimately need to access the Services Director CLI.	The SSH port used by the CLI.	TCP
443	Any machine that may legitimately need to access the Services Director GUI.	The graphical user interface (GUI).	TCP
3306	Services Director HA pair peer.	Used for High Availability operations.	TCP
8100	Any machine that may legitimately need to access the Services Director REST API, including HA pair peer and vTMs using Legacy FLA.	The Services Director REST API. Also used for licensing vTMs that use Legacy FLA Licensing.	TCP
8101	vTMs using Universal FLA.	The Services Director licensing server port. Used for licensing vTMs that use Universal FLA Licensing.	TCP
9070	Services Director HA pair peer.	Used for High Availability operations.	TCP
9080	Services Director HA pair peer.	Used for High Availability operations.	TCP
9090	Services Director HA pair peer.	Used for High Availability operations.	TCP
9091	Services Director HA pair peer.	Used for High Availability operations.	TCP

The following table lists ports must be open on all vTM instances.

Port	Description	Protocol
9070	The REST API port.	TCP
9080	The control port used for cluster operations.	TCP
9090	The graphical user interface (GUI).	TCP
9091	Internal vTM cluster communication.	TCP

Installing the Services Director VA on vSphere

- [Overview: Services Director VA on vSphere](#) 13
- [Obtaining the Services Director VA OVA Package](#) 14
- [Obtaining Services Director Licenses](#) 14
- [Creating a VM in vSphere](#) 14
- [Accessing the Services Director VA on VMware](#) 15

Overview: Services Director VA on vSphere

You can install the Services Director as a Virtual Machine (VM) on the VMware hypervisor.

Perform the following procedure to install and configure the Services Director VA on VMware:

Note: All required files must be in accessible locations in your infrastructure during the installation process. For example, locate the files on an accessible server, or your local machine.

Note: This procedure assumes that you have DHCP or DNS enabled as required by your network.

1. Obtain the Services Director OVA package from Pulse Secure Support. See [“Obtaining the Services Director VA OVA Package” on page 14](#).
2. Obtain the Services Director license from your Pulse Secure account team. For details about obtaining your license keys, see [“Obtaining Services Director Licenses” on page 14](#).
3. Install the Services Director OVA package on vSphere to create the Services Director VA. See [“Creating a VM in vSphere” on page 14](#).
4. Power on the Services Director VA in vSphere and access the Services Director VA with any browser, using its HTTP URL. Log in using the default username (*admin*) and password (*password*).
5. The Setup Wizard runs automatically. Use the wizard to configure your Primary Services Director VA. See [“Running the Services Director VA Setup Wizard” on page 79](#).
6. Review and configure the Settings for the Services Director VA, see [“Installing the Services Director VA on vSphere” on page 13](#).
7. Repeat steps 3 to 5 of this process for the Secondary Services Director to form an HA pair.

Obtaining the Services Director VA OVA Package

The Services Director VA is provided by Pulse Secure Support as an OVA package that contains the VMX and VMDK files necessary to create virtual resources. The Services Director OVA package enables you to create a Services Director VA on ESXi.

You obtain the Services Director OVA package from Pulse Secure support.

Obtaining Services Director Licenses

License tokens are automatically emailed to you when you order your product. If you have not received your license tokens, contact your Pulse Secure sales representative.

Note: If you need assistance locating your local Pulse Secure sales representative, contact Pulse Secure Support.

You must redeem your license tokens at the Pulse Secure License Redemption Portal. To redeem a license token you must have a support site login and password, and a self-signed SSL certificate.

Note: You cannot use a CA-signed certificate.

All licenses are emailed to you as attachments.

Note: You will receive a Legacy FLA License as part of the redemption process. However, if you intend to use only Virtual Traffic Manager (vTM) instances that are at version 10.1 (or later), each with its REST API enabled, you do not need to install this Legacy FLA license. You will instead use a Universal License that comes pre-installed with the Services Director.

Creating a VM in vSphere

To create a virtual machine (VM) in vSphere, you must install the Services Director OVA package on a VMware ESXi host using the vSphere client.

Note: You must be familiar with installing, configuring, and managing VMs using VMware vSphere. The following instructions may vary. For detailed information about creating a VM in vSphere, refer to <http://www.vmware.com/products/vsphere-hypervisor/>.

1. Log in to vSphere.
2. Click **File > Deploy OVF template**. The deployment wizard starts.
3. On the **Source** page, click **Browse**, select the OVA package, click **Open** and then click **Next**.
4. On the **OVF Template Details** page, verify that the OVA package is the one you want to deploy and click **Next**.
5. On the **Name and Location** page, enter a **Name** for the VM and click **Next**.

6. On the **Host/Cluster** page, select a host datastore. This will store the VM and its virtual disk files. Then, click **Next**.

Ensure that the host datastore you select has enough capacity to install the OVA package. See [“Required Hardware Resources for Virtual Machines” on page 10](#).

7. On the **Storage** page, select the required destination storage and a datastore, and click **Next**.
8. On the **Disk Format** page, select the **Thick provisioned** format and click **Next** to pre-allocate all storage.
9. On the **Network Mapping** page, map your *VMNetworkLAN* source network to a destination network using the pull-down list. Then, click **Next**.

There is no need to connect the auxiliary interface. The auxiliary interface can be safely disconnected in the Virtual Machine settings after initial deployment, because this interface is not used by the Pulse Secure Services Director.

10. On the **Ready to Complete** page, verify the deployment settings, select the **Power on after deployment** check box if required, and click **Finish**.
11. When the deployment finishes, click **Close**. The new VM appears under the VM inventory.

You can now configure the Services Director VA using the Setup Wizard, see [“Running the Services Director VA Setup Wizard” on page 79](#).

Accessing the Services Director VA on VMware

To access the Services Director VA, you need the IP address of its management interface.

If DHCP is available, you need to find out the allocated IP address. To do this:

1. Log in to the Services Director VA using the vSphere console.
2. Do not use the jump-start setup wizard.
3. Obtain the allocated DHCP IP address of the VA using the following commands:

```
<host> > enable
<host> # show interfaces
```

If DHCP is *not* available, complete the following steps:

1. Log in to the Services Director VA using the vSphere console.
2. Use the jump-start setup wizard to set:
 - A static IP address.
 - A netmask.
 - The default gateway IP address.

You can access the Services Director VA with a browser, and configure the Services Director VA using the Setup Wizard, see [“Installing the Services Director VA on vSphere” on page 13](#).

Installing the Services Director VA on KVM-QEMU

- [Overview: Services Director VA on KVM-QEMU](#) 17
- [Obtaining the Services Director VA KVM Image](#)..... 18
- [Obtaining Services Director Licenses](#)..... 18
- [Creating the Services Director VA on a KVM Server](#) 18
- [Accessing the Services Director VA on KVM](#) 25

Overview: Services Director VA on KVM-QEMU

The Pulse Secure Services Director Virtual Appliance is supported for production use on the KVM-QEMU hypervisor running on either an Ubuntu 18.04 or a RHEL/CentOS 6.x/7.x server.

Note: The Services Director VA is available on KVM-QEMU as a 64-bit version only.

Perform the following steps to install and configure the Services Director VA on KVM-QEMU:

Note: All required files must be in accessible locations in your infrastructure during the installation process. For example, locate the files on an accessible server, or your local machine.

Note: This procedure assumes that you have DHCP or DNS enabled as required by your network.

1. Obtain the Services Director Kernel Virtual Machine (KVM) image in QCOW2 format from Pulse Secure Support. See [“Obtaining the Services Director VA KVM Image” on page 18](#).
2. Obtain the Services Director license from your Pulse Secure account team. For details about obtaining your license keys, see [“Obtaining Services Director Licenses” on page 18](#).
3. Prepare a server that supports KVM. Supported servers are Ubuntu 18.04 and RHEL/CentOS 6.x/7.x.
4. Install the Services Director QCOW2 virtual machine on your server. This process creates the Services Director VA. See [“Creating the Services Director VA on a KVM Server” on page 18](#).
5. Access the Services Director VA. See [“Accessing the Services Director VA on KVM” on page 25](#).
6. The Setup Wizard runs automatically. Use the wizard to configure your Primary Services Director VA. See [“Running the Services Director VA Setup Wizard” on page 79](#).
7. Review and configure the Settings for the Services Director VA, see [“Installing the Services Director VA on KVM-QEMU” on page 17](#).
8. Repeat steps 3 - 6 for the Secondary Services Director to form a High Availability (HA).

Obtaining the Services Director VA KVM Image

The Services Director VA is provided by Pulse Secure Support as a KVM image in QCOW2 format. This image contains the files necessary to create a Services Director VA on a KVM-QEMU hypervisor on all supported server platforms.

You obtain the Services Director KVM image in QCOW2 format from Pulse Secure Support.

Obtaining Services Director Licenses

License tokens are automatically emailed to you when you order your product. If you have not received your license tokens, contact your Pulse Secure sales representative.

You must redeem your license tokens at the Pulse Secure License Redemption Portal. To redeem a license token you must have a support site login and password, and a self-signed SSL certificate.

Note: You cannot use a CA-signed certificate.

All licenses are emailed to you as attachments.

Note: You will receive a Legacy FLA License as part of the redemption process. However, if you intend to use only Virtual Traffic Manager (vTM) instances that are at version 10.1 (or later), each with its REST API enabled, you do not need to install this Legacy FLA license. You will instead use a Universal License that comes pre-installed with the Services Director.

Creating the Services Director VA on a KVM Server

To create the Services Director VA on a KVM server, you must install the Services Director KVM image on a KVM server. There are many virtualization systems in common use; the following two examples describe the installation of your Services Director VA:

- Using the command line interface (CLI) of the libvirt toolset. See **“Creating a VM Using the libvirt Command Line Interface” on page 19.**

For detailed information about libvirt, refer to <https://libvirt.org/>.

- Using the graphical user interface (GUI) of the Virtual Machine Manager graphical toolset. See **“Creating a VM Using the VMM Graphical User Interface” on page 20.**

For detailed information about VMM, refer to <https://virt-manager.org/>.

However your image is installed, the following settings must be used for the virtual machine:

- X86_64 architecture.
- Four virtual CPUs.
- 8192 MB (8 GB) of memory.
- Write-through caching mode.
- Two Ethernet adapters with an e1000 model, connected using a bridge.
- A hard drive with IDE or VIRTIO bus type for the KVM image in QCOW2 format.

Note: The installation and configuration of your chosen toolset is outside the scope of this document. Refer to your tool's documentation for details.

Creating a VM Using the libvirt Command Line Interface

To perform this procedure, you must have the required tools installed on a KVM-QEMU hypervisor, and be familiar with installing, configuring, and managing VMs.

1. Copy the KVM image to an appropriate designated directory (storage pool). Your System Administrator determines which storage pool to use. Give the file a unique name. For example, the filename might be of the form "image_xx.qcow2". Images can only be used once.

For the purposes of this example, this directory is `/vms/pool/sd0`.

2. Install the required VM by issuing a `virt-install` command using the following syntax:

```
virt-install --import
--name=<servicedirector_name>
--disk <image_pool_path>/image.qcow2,format=qcow2,bus=<bus>,cache=writethrough
--os-type=linux
--network bridge=<bridge_name>,model=<model for primary interface>
--network bridge=<bridge_name>,model=<model for auxiliary interface>
--ram=8192 --arch=x86_64 --vcpus=4
```

Where bus can be set to either 'ide' or 'virtio'.

For example:

```
virt-install --import
--name=sd_kvm_07
--disk /vms/pool/sd0/image.qcow2,format=qcow2,bus=ide,cache=writethrough
--os-type=linux
--network bridge=br0,model=e1000
--network bridge=br0,model=e1000
--ram=8192 --arch=x86_64 --vcpus=4
```

After the installation completes, a number of background initialization tasks take place. As a result, the CLI will offer reduced functionality for a short period. Pulse Secure recommends waiting at least two minutes before attempting to access the Services Director.

3. List the VMs on this hypervisor:

```
virsh list
```

The response includes your VM (along with other VMs, if any):

Id	Name	State
356	pchaudh-07	running
542	sramakrishnan-0b	running
593	sd_kvm_07	running

4. Access the console of the VM you have just deployed:

```
virsh console <vm_name>
```

For example:

```
virsh console sd_kvm1
```

To exit the console, use ctrl+].

Creating a VM Using the VMM Graphical User Interface

To perform this procedure, you must have the required tools installed on a KVM-QEMU hypervisor, and be familiar with installing, configuring, and managing VMs.

1. Copy the KVM image to an appropriate designated directory (storage pool). Your System Administrator determines which storage pool to use. The image filename must be "image.qcow2".

For the purposes of this example, this directory is */var/lib/libvirt/images*.

2. Start the VMM GUI:

```
virt-manager --connect=qemu+ssh://my-kvm-host.com/system
```

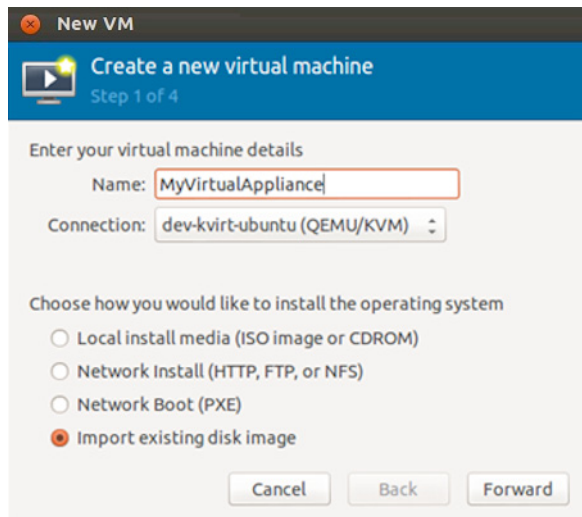
In this command, *my-kvm-host.com* is the host machine name.

An SSH tunnel is used to connect to the KVM-QEMU host. You must have an SSH account and corresponding public key stored on this machine for authentication.

Refer to the VMM documentation for information on alternative connection methods.

- Click **New** to start the process of creating a new virtual machine.


FIGURE 1 Creating a New Virtual Machine Wizard: 1 of 4



- Enter a **Name** for your virtual appliance that corresponds with the name used for the disk image file.
- Select **Import existing disk image** from the list of options.
- Click **Forward** to proceed.

The next page of the wizard appears:


FIGURE 2 Creating a New Virtual Machine Wizard: 2 of 4



- Click **Browse** to select the storage pool location and disk image file for this virtual machine.
- Ensure that the **OS type** is Generic.
- Ensure that the **Version** is Generic.

10. Click **Forward** to proceed. The next page of the wizard appears:

FIGURE 3 Creating a New Virtual Machine Wizard: 3 of 4

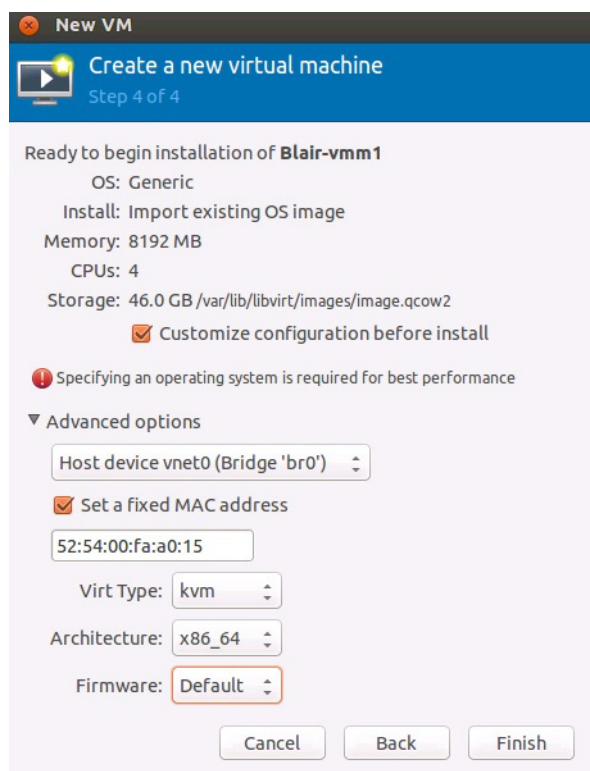


11. Set the **Memory (RAM)** to 8192 MB

12. Set the **CPUs** to 4.

13. Click **Forward** to proceed. The next page of the wizard appears:

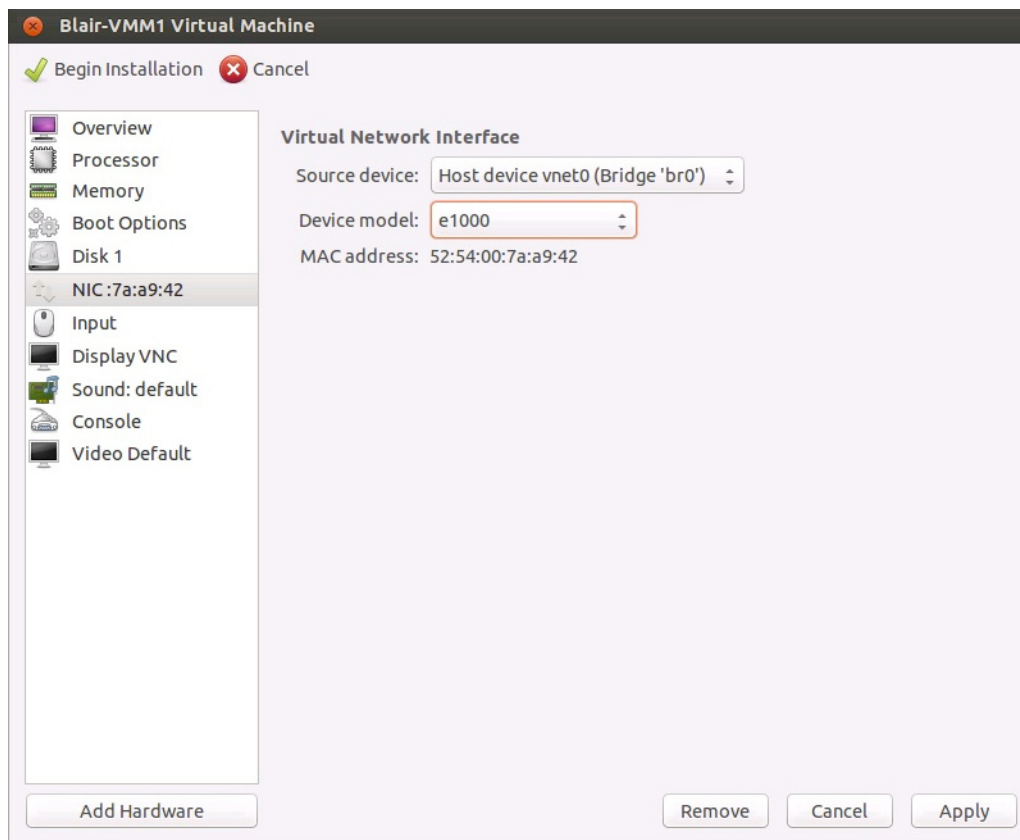
FIGURE 4 Creating a New Virtual Machine Wizard: 4 of 4



14. Check that the summary information is correct.

15. Ensure that the **Customize configuration before install** check box is selected.
16. Expand **Advanced options**.
17. Set **Architecture** to x86_64.
18. Click **Finish**. A configuration dialog box appears.
19. Select **Disk 1** to update disk settings:
 - Under **Advanced Options**, ensure that **Storage format** is set to qcow2.
 - Under **Advanced Options**, ensure that **Disk bus** is set to either IDE or Virtio.
 - Under **Performance Options**, ensure that **Cache mode** is set to writethrough.
 - Click **Apply**.
20. Select **Virtual Network Interface** to view Virtual Network Interface settings.

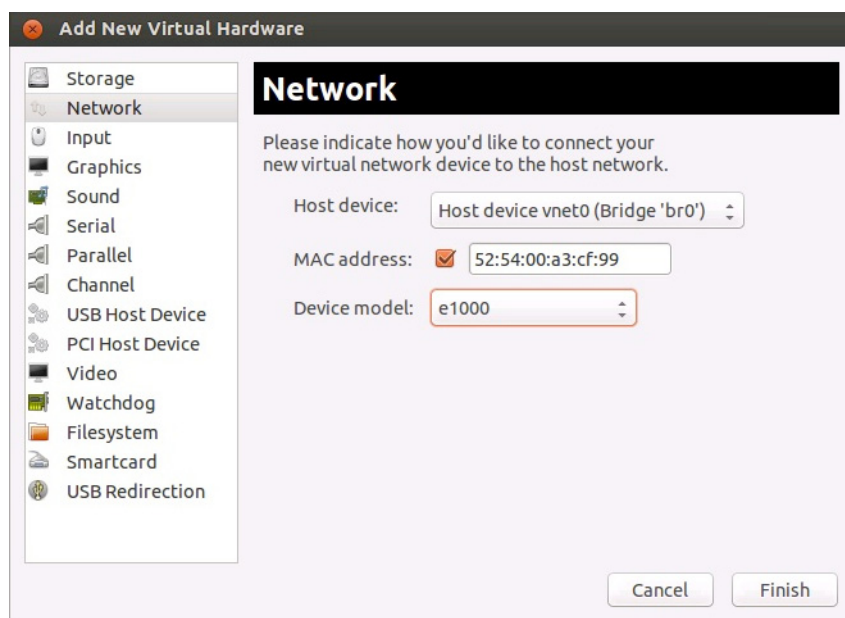
FIGURE 5 Configuring the KVM Virtual Machine: Virtual Network Interface



21. Ensure that the **Source device** is the br0 bridge.
22. Set the **Device model** to e1000.
23. Click **Apply**.

24. Click **Add Hardware**.
25. Click **Network**. The dialog box updates.

FIGURE 6 Configuring the KVM Virtual Machine: Network



26. Ensure that the **Host device** is the br0 bridge.
27. Set the **Device model** to e1000.
28. Click **Finish**.
29. Select **Begin installation** to complete the installation process.

After the installation completes, a number of background initialization tasks take place. As a result, the CLI will offer reduced functionality for a short period. Pulse Secure recommends waiting at least two minutes before attempting to access the Services Director VA.

Accessing the Services Director VA on KVM

To access the Services Director VA, you need the IP address of its management interface.

If DHCP is available, you need to find out the allocated IP address.

1. Log in to the Services Director VA using the KVM console.

Do not use the jump-start setup wizard.

2. Obtain the allocated DHCP IP address of the VA using the following commands:

```
<host> > enable  
<host> # show interfaces
```

If DHCP is *not* available, complete the following steps:

1. Log in to the Services Director VA using the KVM console.
2. Use the jump-start setup wizard to set:
 - A static IP address.
 - A netmask.
 - The default gateway IP address.

You can access the Services Director VA with a browser, and configure the Services Director VA using the Setup Wizard, see [“Installing the Services Director VA on KVM-QEMU” on page 17](#).

Installing the Services Director VA on Amazon Web Services

- **Overview: Services Director VA on Amazon Web Services** 27
- **Obtaining Services Director Licenses** 27
- **Launching and Configuring the Primary Services Director on AWS** 28
- **Launching and Configuring the Secondary Services Director on AWS** 78

Overview: Services Director VA on Amazon Web Services

Services Director instances can be launched from Amazon Machine Images (AMIs) on Amazon Web Services (AWS). AWS supports HA pairing of two Services Director nodes. Each Services Director is launched from a separate AMI and then configured and joined.

To create an HA pair of Services Director nodes on AWS:

1. Obtain the Services Director license from your Pulse Secure account team. For details about obtaining your license keys, see **“Obtaining Services Director Licenses” on page 27**.
2. Create the Primary Services Director node, see **“Launching and Configuring the Primary Services Director on AWS” on page 28**.
3. Create the Secondary Services Director node, see **“Launching and Configuring the Secondary Services Director on AWS” on page 78**.
4. Review and configure the Settings for the Services Director VA, see **“Installing the Services Director VA on Amazon Web Services” on page 27**.

Obtaining Services Director Licenses

License tokens are automatically emailed to you when you order your product. If you have not received your license tokens, contact your Pulse Secure sales representative.

You must redeem your license tokens at the Pulse Secure License Redemption Portal. To redeem a license token you must have a support site login and password, and a self-signed SSL certificate.

Note: You cannot use a CA-signed certificate.

All licenses are emailed to you as attachments.

Launching and Configuring the Primary Services Director on AWS

Perform the following procedure to launch and configure a Primary Services Director VA on AWS:

1. Prepare the required infrastructure (VPCs and Subnets) in the AWS network, see [“Preparing AWS Infrastructure” on page 28](#).
2. Prepare an AWS Security Group, see [“Preparing an AWS Security Group” on page 40](#).
3. Launch a Services Director instance on AWS from the Services Director AMI, see [“Launching a Services Director AMI Instance on AWS” on page 64](#).
4. (Optional) Add and configure elastic IP addresses for the Primary Services Director instance, see [“Creating Elastic IP Addresses for the Services Director Instance” on page 69](#).
5. Update your AWS Security Group to include all allocated IP addresses for the Primary Services Director instance, see [“Updating Security Rules for Services Director Instance IP Addresses” on page 74](#).
6. Retrieve the password for the Primary Services Director from AWS, see [“Retrieving the Default Password for a Services Director Instance” on page 74](#).
7. Access the Primary Services Director instance using a browser, see [“Accessing your Services Director Instance for the First Time” on page 77](#).
8. Use the Setup Wizard (which starts automatically) to create your Primary Services Director node. See [“Running the Services Director VA Setup Wizard” on page 79](#).

Preparing AWS Infrastructure

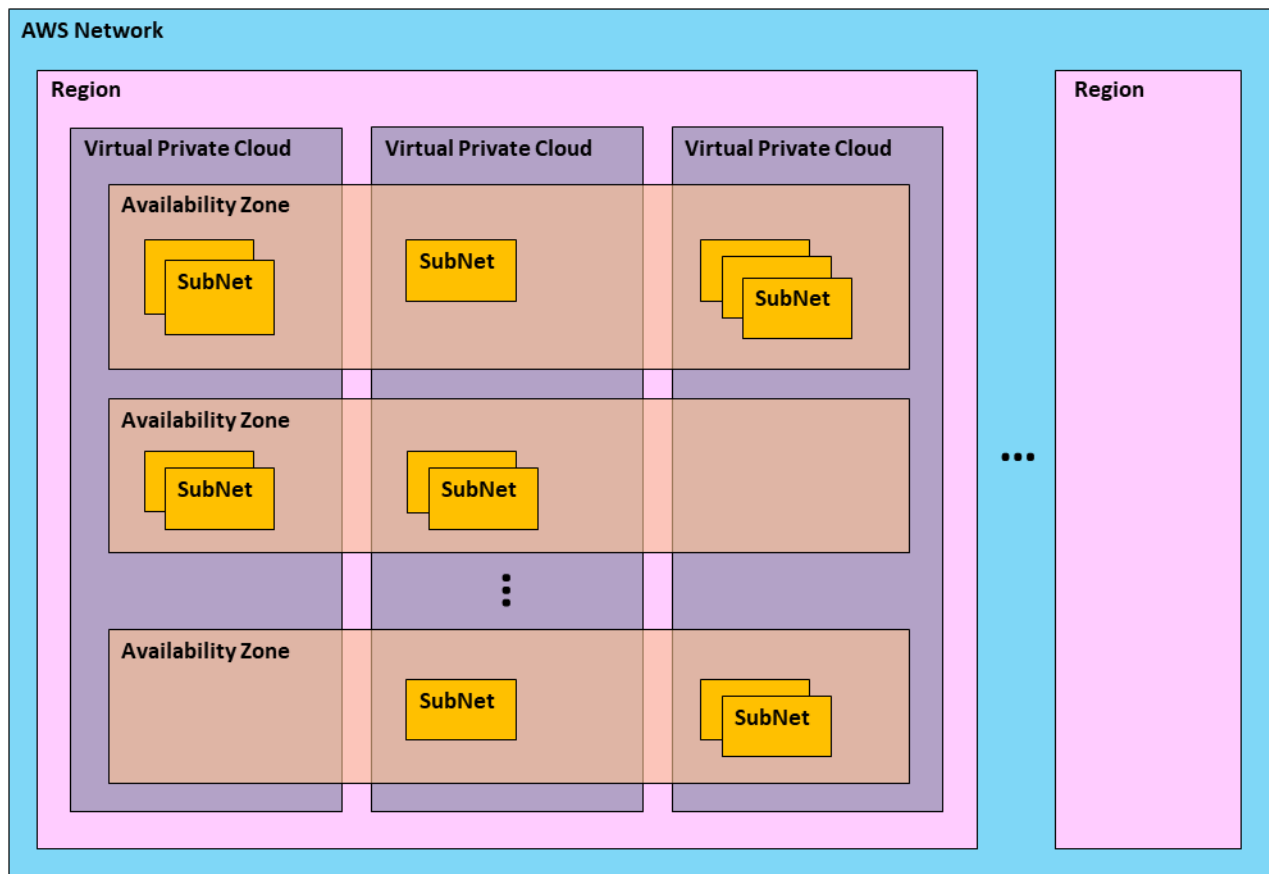
Before you can launch a pair of Services Director VA nodes into AWS, you must prepare any required AWS infrastructure elements within the AWS Network. This requires:

- [“Understanding AWS Infrastructure” on page 29](#).
- [“Determining IP Address Requirements” on page 30](#).
- [“Creating an AWS Virtual Private Cloud” on page 35](#).
- [“Creating AWS Subnets” on page 37](#).

Understanding AWS Infrastructure

The following diagram shows AWS infrastructure concepts and relationships.

FIGURE 7 AWS Networking Infrastructure



The AWS infrastructure and the relationships between each type are as follows:

- The *AWS Network* is a secure cloud services platform.
The AWS Network has many AWS Regions.
- A *Region* is a named set of AWS resources based in the same geographical area, such as a country.
Every Region has at least two AWS Availability Zones.
- An *Availability Zone* is a geographical location entirely within a Region. The geographic nature of an Availability Zone insulates it from service failures in other Availability Zones.
Each Availability Zone supports network access to all other Availability Zones in the Region.
Each Availability Zone is accessible by every AWS Virtual Private Cloud in a Region.
- A *Virtual Private Cloud (VPC)* is a virtual network. It is populated by AWS infrastructure elements that share network security and connectivity.

A VPC can access all Availability Zones in the Region.

A VPC requires one or more AWS Subnets.

VPCs are created and managed by the customer.

The required VPC (and its Subnets) must be in place before a Services Director pair can be launched, see **“Creating an AWS Virtual Private Cloud” on page 35.**

- A *Subnet* is a subdivision of the IP address range of a VPC.

Subnets are created by the customer to group application instances according to security and operational needs. A Subnet is entirely contained within a single Availability Zone. Each Services Director is launched into a Subnet. Any required Subnet(s) must be in place before Services Director can be launched, see **“Creating AWS Subnets” on page 37.**

Note: Additional Subnets may also be required for vTM instances.

Determining IP Address Requirements

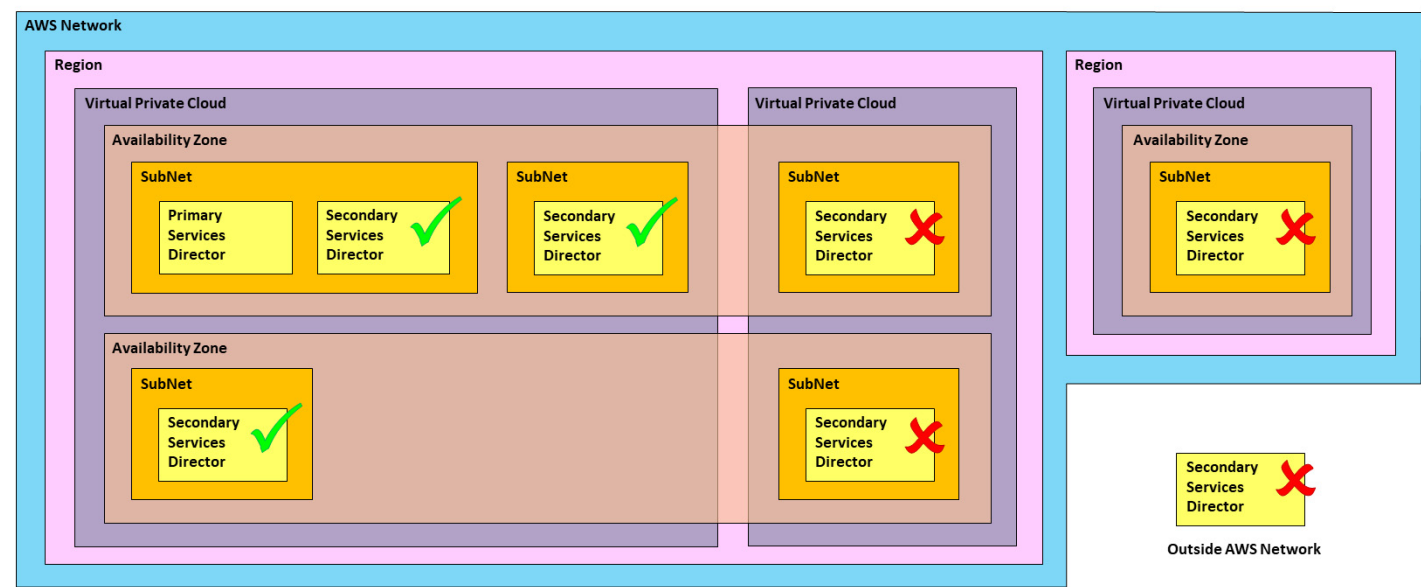
The use of IP address types on AWS (*private*, *elastic* and *public*) are determined by your general networking requirements, but in Services Director terms the following contribute to this choice:

- The relative placement of the Services Director nodes.
- The placement of vTMs relative to the Services Directors.
- Your access requirements for your individual Services Director nodes.

A Secondary Services Director node *must be in the same VPC* as the Primary node. That is:

- In the same Subnet as the Primary node, OR
- In a different Subnet to the Primary node, but in the same Availability Zone as the Primary node's Subnet, OR
- In a different Subnet to the Primary node, but in a different Availability Zone to the Primary node's Subnet.

FIGURE 8 Supported Placement of Secondary Services Director

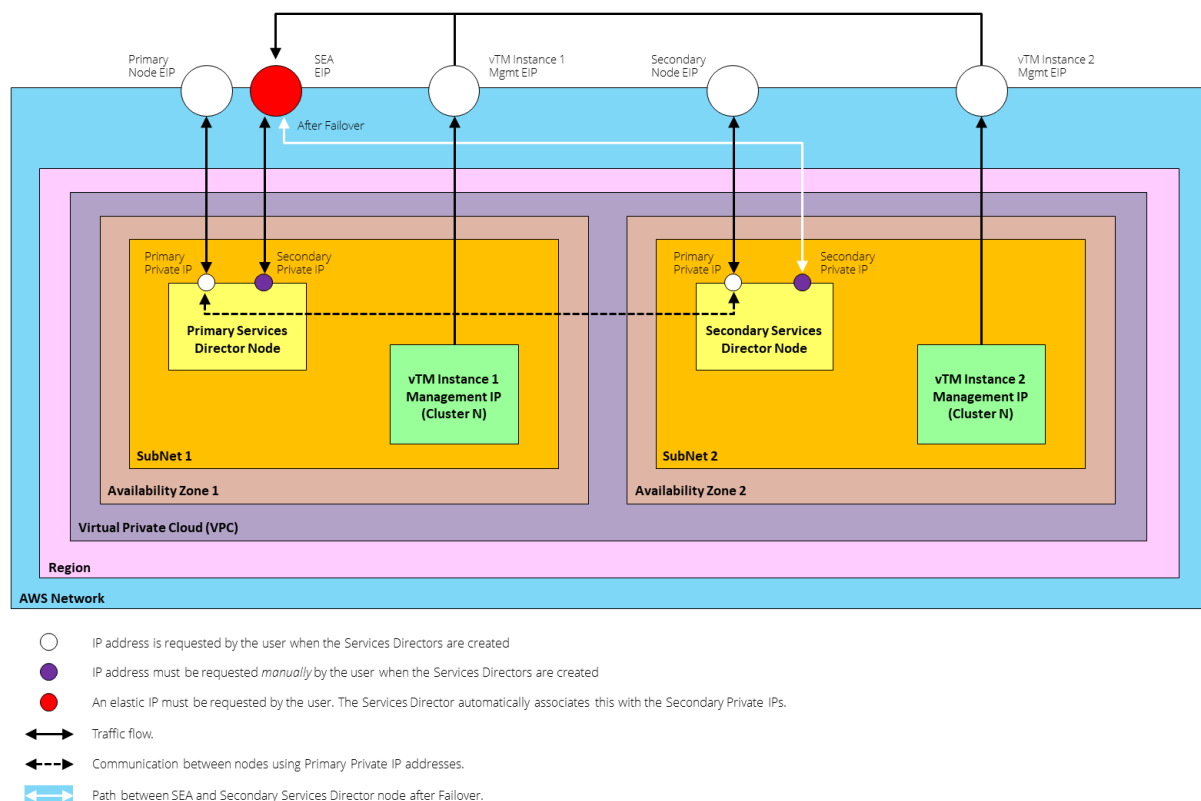


All other Services Director placements (shown above) are not supported.

The specifics of an AWS deployment determine whether private IP addresses or elastic IP addresses are used:

- The use of an Elastic Service Endpoint Address (SEA) is mandatory where the Primary and Secondary Services Directors are in different Subnets. For example, when the vTMs are in the same Availability Zone (or Subnet) of the VPC:

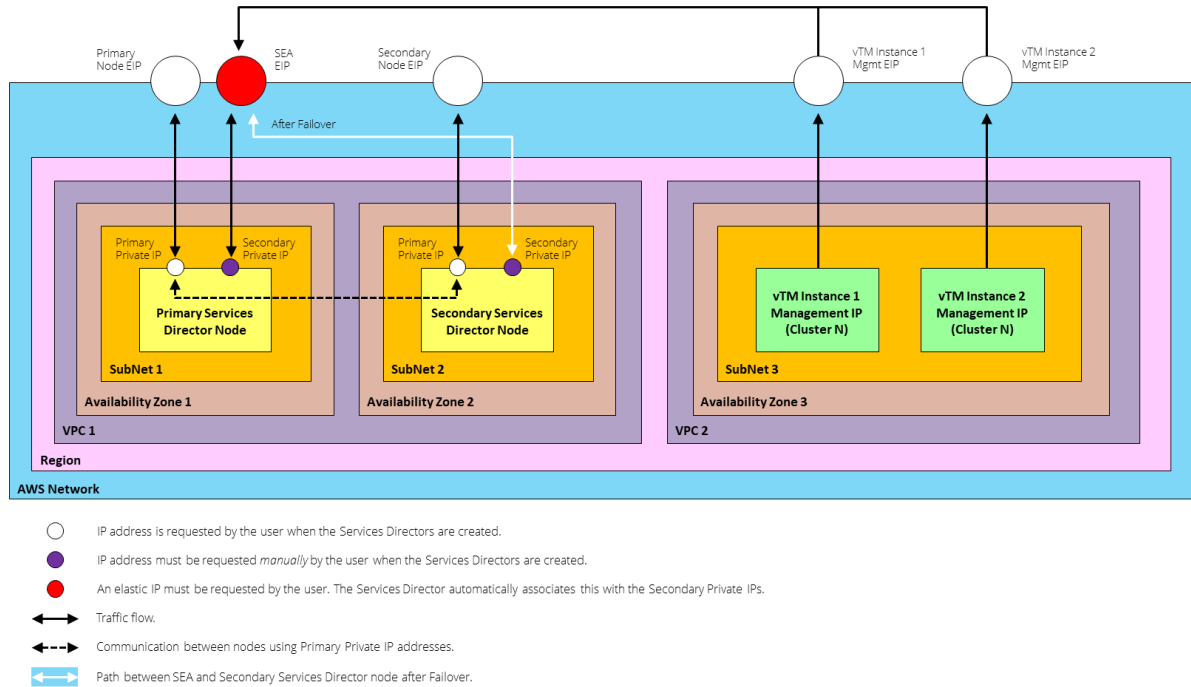
FIGURE 9 Elastic IP Addresses: vTMs in the Same VPC



Note: The placement of vTMs in this example is illustrative; each vTM can be in any Subnet within the VPC containing the Services Director nodes.

Alternatively, when the vTMs are in a different VPC:

FIGURE 10 Elastic IP Addresses: vTMs in Different VPCs



Note: The placement of vTMs in this example is illustrative; they can be in any Subnet outside the Services Directors VPC, or outside AWS completely.

In both cases:

- The *Primary Private IP* address of each Services Director node is used for inter-node communications. Typically, this communication is direct within the VPC, but for inter-node database access and replication the communication will normally be directed via the external SEA.
- The SEA directs traffic to the *Active* Services Director node at all times, using the *Secondary Private IP* Address of the node. Each Services Director node requires an additional Secondary Private IP Address. You must request that AWS auto-assigns an additional Secondary Private IP Address during creation of the Services Director instance.
- You will typically configure Elastic IP addresses for each individual Services Director node, and then associate the Elastic IP for each node with that node's Primary Private IP address.
- Additionally, each vTM will typically have an Elastic IP address that communicates with the Services Director via the SEA.

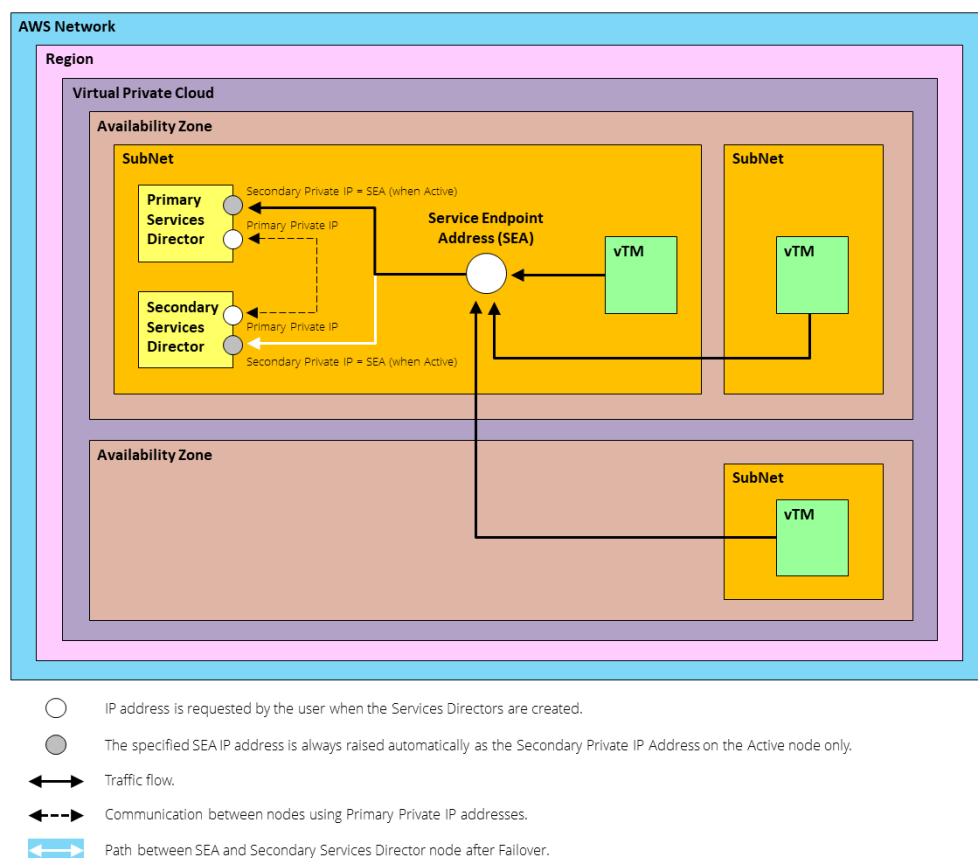
Note: An elastic SEA may also be used when the Primary and Secondary SDs are in the *same* Subnet, but whether this is necessary depends on your external connectivity requirements. For example, when your VPCs are external to the VPC, and will need to access the Services Director for licensing.

Note: Once an Elastic IP address for the SEA is defined, you must configure your network accordingly to ensure connectivity of all components. *This is outside the scope of this document, see the Virtual Traffic Manager documentation.*

Note: The placement of vTMs is only restricted by your networking configuration. All vTMs must connect to the Services Director using the SEA. *This is outside the scope of this document.*

- Where the Services Director pair are in a single Subnet, and all vTMs in its estate are inside a single VPC, you can use Private IP addresses for the Services Director nodes and for the SEA. For example:

FIGURE 11 Private IP Addresses



The *Primary Private IP* address of each Services Director node is used for inter-node communications. Typically, this communication is direct within the VPC, but for inter-node database access and replication the communication will normally be directed via the SEA.

The SEA directs traffic to the *Active* Services Director node at all times. This uses a *Secondary Private IP* address that is raised on the *Active* node only. When failover occurs, this IP address is removed from one node and raised on the other node.

Note: This is standard Services Director behaviour. *You do not have to raise a Secondary Private IP address on either node.*

The vTMs inside the VPC connect (call back) to the Services Directors using the Private IP Address SEA.

If you want to access your individual Services Director nodes from outside the VPC, you can allocate a Public IP address or Elastic IP address to each node.

All IP addresses are required for the definition of an AWS Security Group, see [“Preparing an AWS Security Group” on page 40](#).

Once you have prepared the required VPC and Subnet(s), you can launch a Services Director VA using an AMI on AWS, see [“Launching a Services Director AMI Instance on AWS” on page 64](#).

Creating an AWS Virtual Private Cloud

The required AWS Virtual Private Cloud (VPC) must exist before you can launch the Services Director VA on AWS. You create a VPC from the Amazon Web Services platform.

There can be several VPCs within an AWS Region.

The IP range of a VPC is defined by a CIDR block. For example, *10.0.0.0/16*.


The VPC will contain both nodes of a Services Director HA pair, though the two nodes can be in different Availability Zones within the VPC, see [“Determining IP Address Requirements” on page 30](#).

A VPC can access all AWS Availability Zones in an AWS Region.

To create a VPC:

1. Login to the *AWS Management Console*.
2. On the top bar of the AWS Management Console, select the required **Region**. For example, *EU (Ireland)*.

FIGURE 12 Select AWS Region

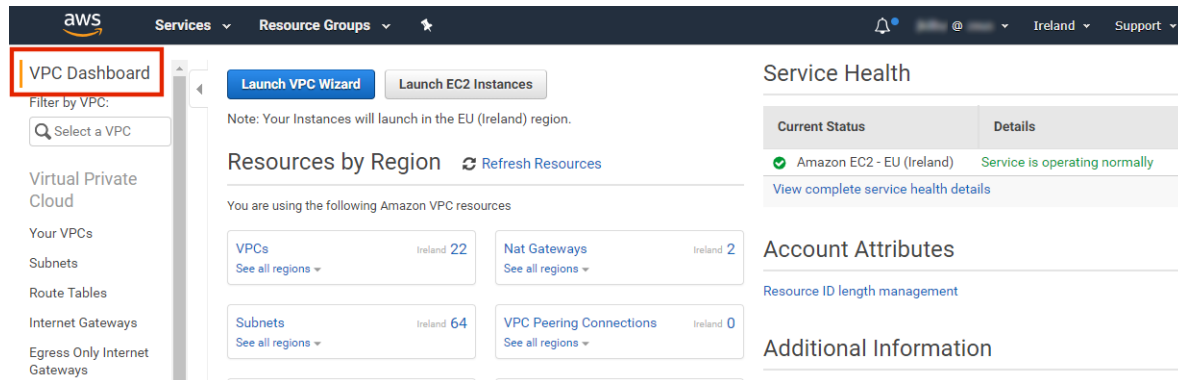


3. On the AWS top bar, click **Services** and then locate the **Network & Content Delivery** options.

- Under **Network & Content Delivery**, select **VPC**.

The **VPC Dashboard** appears. For example:

FIGURE 13 AWS VPC Dashboard

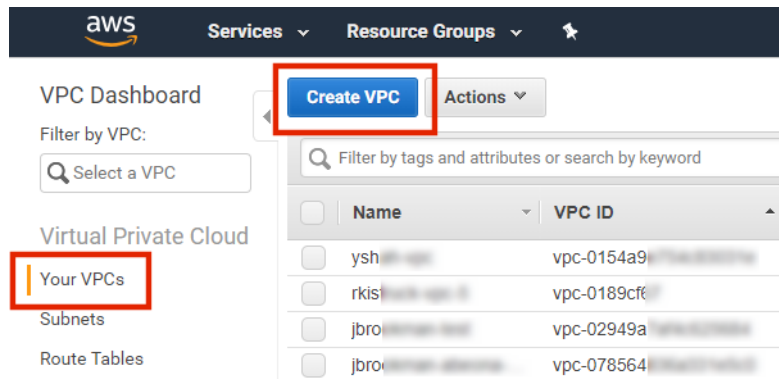


- Under **Virtual Private Cloud**, click **Your VPCs**.

A list of your existing VPCs appears.

- Examine your VPCs and decide if an existing one matches your networking requirements for your Services Director. If there is a suitable VPC, no further actions is required, and this process is complete.
- Click **Create VPC**.

FIGURE 14 Create VPC



The **Create VPC** page appears.

FIGURE 15 Create VPC Properties

VPCs > Create VPC

Create VPC

A VPC is an isolated portion of the AWS cloud populated by AWS objects, such as Amazon EC2 instances. You must specify an IPv4 address range for your VPC. Specify the IPv4 address range as a Classless Inter-Domain Routing (CIDR) block; for example, 10.0.0.0/16. You cannot specify an IPv4 CIDR block larger than /16. You can optionally associate an Amazon-provided IPv6 CIDR block with the VPC.

Name tag ⓘ

IPv4 CIDR block* ⓘ

IPv6 CIDR block ☒ No IPv6 CIDR Block ⓘ ☐ Amazon provided IPv6 CIDR block

Tenancy ⓘ

* Required Cancel Create

- Specify your required networking details and click **Create**.

A confirmation message appears.

- Click **Close**.

The new VPC is added to the **Your VPCs** list.

Once you have a suitable VPC, you can create any required Subnets inside the VPC, see [“Creating AWS Subnets” on page 37](#).

Creating AWS Subnets

The required AWS Subnets must exist inside your chosen VPC before you can launch the Services Director VA on AWS. You create Subnets from the Amazon Web Services platform.

Each Subnet has an IP address range that is a subdivision of the VPC's total range. The range is that is defined by a CIDR block. For example, 10.0.0.0/24.

The Subnet will contain either one or both of the Services Director nodes, see [“Determining IP Address Requirements” on page 30](#).

The Subnet can be inside any AWS Availability Zone within the VPC.

To create a Subnet:

- Login to the *AWS Management Console*.
- On the top bar of the AWS Management Console, select the required **Region**.
- On the AWS top bar, click **Services** and then locate the **Network & Content Delivery** options.

- Under **Network & Content Delivery**, select **VPC**.

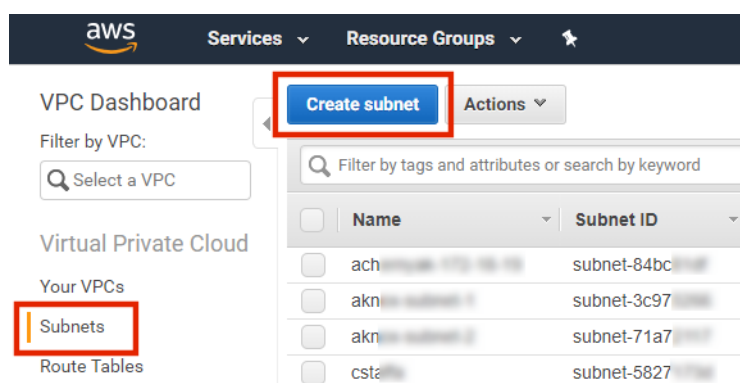
The **VPC Dashboard** appears.

- Under **Virtual Private Cloud**, click **Subnets**.

A list of your existing Subnets appears.

- Examine your Subnets and decide if you have a Subnet(s) that match your networking requirements for your Services Director nodes. If there is a suitable Subnet(s), no further actions is required, and this process is complete.
- Click **Create subnet**.

FIGURE 16 Create Subnet



The **Create subnet** page appears.

FIGURE 17 Create Subnet Properties

[Subnets](#) > Create subnet

Create subnet

Specify your subnet's IP address block in CIDR format; for example, 10.0.0.0/24. IPv4 block sizes must be between a /16 netmask and /28 netmask, and can be the same size as your VPC. An IPv6 CIDR block must be a /64 CIDR block.

Name tag ⓘ

VPC* ⓘ

VPC CIDRs	CIDR	Status	Status Reason

Availability Zone ⓘ

IPv4 CIDR block* ⓘ

* Required

[Cancel](#) [Create](#)

8. Specify your required networking details and click **Create**.

A confirmation message appears.

9. Click **Close**.

The new Subnet is added to the **Subnets** list.

10. Repeat steps 7 - 9 if a second Subnet is required for your Secondary Services Director node.

Note: The relative positions of your Subnets for the Services Director nodes will influence your choice of IP address, see [“Determining IP Address Requirements” on page 30](#).

Once you have a suitable Subnet(s), you can prepare your AWS Security Group, see [“Preparing an AWS Security Group” on page 40](#).

Preparing an AWS Security Group

Before you launch a Services Director from an AMI on AWS, you must define an *AWS Security Group*.

An AWS Security Group is a named set of permitted inbound network connections that apply to one or more AWS AMI instances. Each Security Group consists of a list of rules. Each rule identifies a protocol, a port, and an IP address (or IP address range) from which inbound requests can be received.

Note: Rules can reference the Security Group itself. This represents all traffic from any IP address hosted on a AWS instance that uses the Security Group.

All received requests that are not permitted by a rule are refused.

In Services Director terms, each Services Director node is an AWS AMI instance, and the assigned Security Group will control which inbound connections can reach the Services Director node via its SEA.

Note: Security Groups also optionally support rules to govern permitted outbound connections. This guide does not specify suitable rules to govern permitted outbound connections from Services Director.

There are three general deployment scenarios for your Services Directors and vTMs. Your chosen scenario determines the required rules for your Security Group:

- Using elastic IP addresses for the Services Director nodes and the SEA, with the vTMs in the same VPC as the Services Directors, see **“Scenario 1 - Elastic IPs with vTMs in the Same VPC” on page 41**.
- Using elastic IP addresses for the Services Director nodes and the SEA, with the vTMs in a different VPC or outside AWS completely, see **“Scenario 2 - Elastic IPs with vTMs in Different VPCs” on page 47**.
- Using private IPs for the Services Director nodes, see **“Scenario 3 - Private IPs Only” on page 53**.

Once you have identified your network configuration and the required rules, you can create your AWS Security Group. see **“Creating an AWS Security Group” on page 59**.

After you launch the AMI for each Services Director node, you must create additional rules:

- In your Security Group used by the Services Directors, you add Security Group rules for the IP addresses used by each Services Director. This enables the Services Directors to communicate with each other **“Updating Security Rules for Services Director Instance IP Addresses” on page 74**.
- In any Security Groups used by the vTMs that will be in the estate of the Services Director, you add Security Group rules for the Internet-facing IP SEA. This enables the vTMs to communicate with the Services Director.

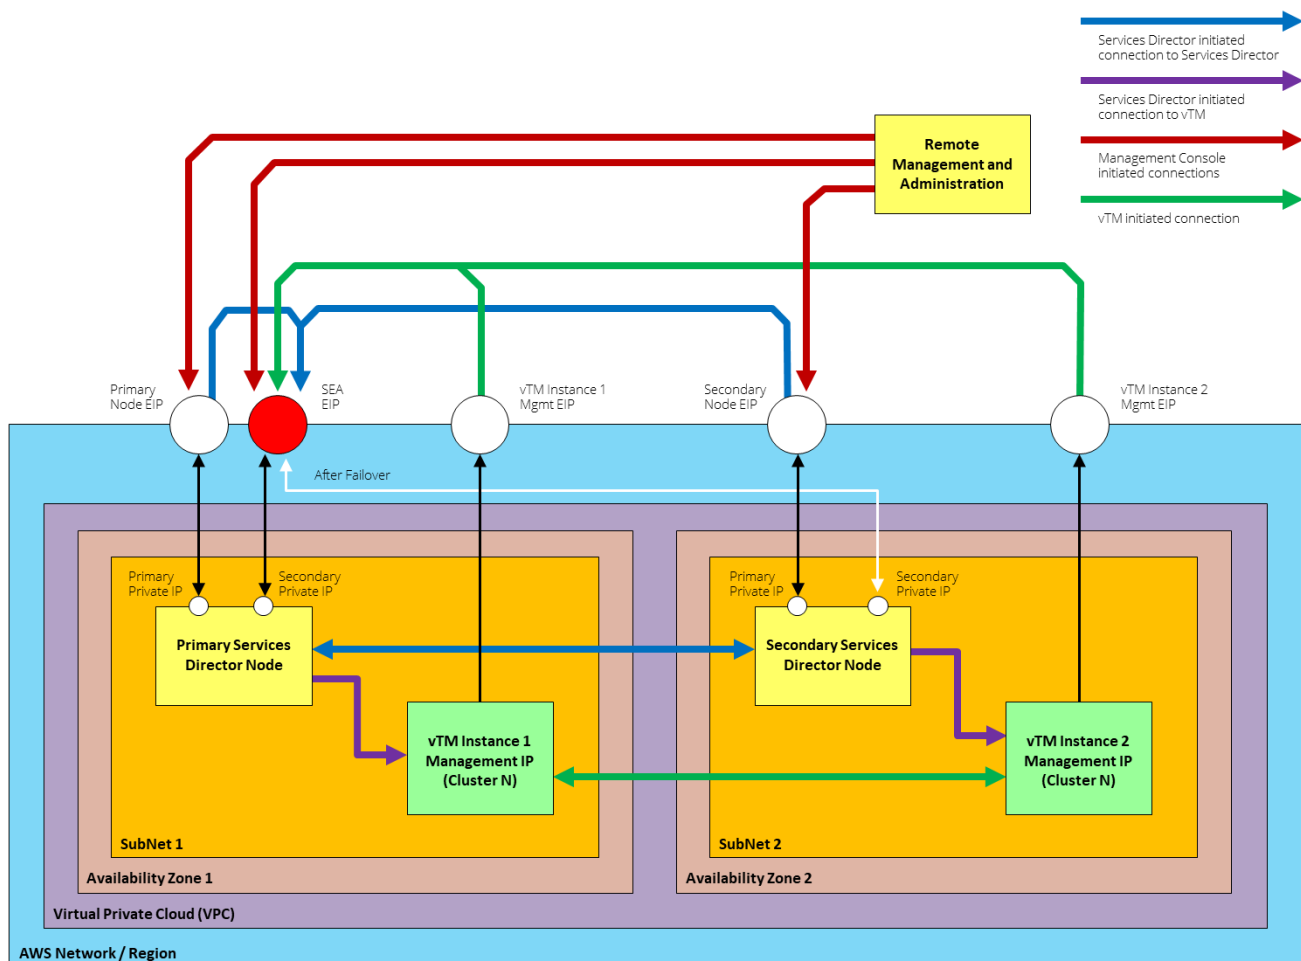
Scenario 1 - Elastic IPs with vTMs in the Same VPC

In this scenario:

- AWS Elastic IPs (EIPs) are used for the Primary and Secondary nodes, and for the Services Endpoint Address (SEA).
- The Services Director nodes are in the same AWS VPC, although they can be in separate Subnets or availability zones.
- The vTMs in the estate of the Services Director pair are in the same VPC.
- The self-registration feature of vTM is supported when the vTM is provided with the SEA EIP of the Services Director.

The flow of requests through the defined IP addresses is as follows:

FIGURE 18 Elastic IPs with vTMs in the Same VPC



Note: When vTM Communications Channel (Comms Channel) is in use, there are minor differences in the scenario diagram above. All Services Director-initiated connections to the vTM (purple in the diagram) are replaced by vTM-initiated connections to the Services Director (green in the diagram). See [“Using Scenario 1 with vTM Communications Channel”](#) on page 42.

Note: When you join a Secondary Services Director to a Primary Services Director, always specify the Primary Private IP Address of the Primary Services Director. This allows the majority of traffic between the Services Director nodes to be routed within the VPC and not over the public Internet.

To set up the required AWS Security Group for this scenario, you must create the rules listed in the following sections:

- **“Services Director Security Group: Remote Management and Administration” on page 43.**
- **“Services Director Security Group: Services Director Peer” on page 43.**
- **“Services Director Security Group: vTM Estate” on page 44.**
- **“vTM Security Group: Remote Management and Administration” on page 45.**
- **“vTM Security Group: vTM Peers” on page 45.**
- **“vTM Security Group: Services Director Estate Manager” on page 46.**

Note: For all rules, CIDR format must be used for each IP address range. This takes the form xx.xx.xx.xx/nn, where nn is the size of the subnet mask. The subnet mask size cannot be omitted, even when you specify a single IP address. That is, if you require 10.11.12.13 as an IP address, you must specify 10.11.12.13/32.

Using Scenario 1 with vTM Communications Channel

When vTM Communications Channel (Comms Channel) is in use, there are minor differences from the scenario described above. All Services Director-initiated connections to the vTM (purple in the diagram) are replaced by vTM-initiated connections to the Services Director (green in the diagram).

If you are using Comms Channel, you should not attempt to connect to the Services Director from the vTM until you have its Elastic IP address. That is, you must perform the following tasks in order:

1. Create the vTM in AWS. **“Launching and Configuring the Primary Services Director on AWS” on page 28.**
2. Run the Configuration Wizard for the vTM, but do not specify Services Director details for the vTM.
3. Associate an Elastic IP address to the vTM’s Private IP address. **“Creating Elastic IP Addresses for the Services Director Instance” on page 69.**
4. Configure Security Groups on the Services Director, using the Elastic IP for the vTM (refer to the sections that follow).
5. Log into the vTM and self-register the vTM with the Services Director, see **“Requesting Self Registration on a Configured vTM” on page 177.**

Services Director Security Group: Remote Management and Administration

You must add the following rules to the Security Group *used by the Services Director*.

The following rules support flows from Remote Management and Administration.

In these rules, the IP address of the device from which you will perform remote management is required.

Note: For all rules, CIDR format must be used for any IP address range. This takes the form xx.xx.xx.xx/nn, where nn is the size of the subnet mask. The subnet mask size cannot be omitted, even when you specify a single IP address. That is, if you require 10.11.12.13 as an IP address, you must specify 10.11.12.13/32.

Type	Protocol	Port Range	Source	Description
SSH	TCP	22	Custom Remote Management IP Address.	Administrative shell access to Services Director.
HTTPS	TCP	443	Custom Remote Management IP Address.	Administrative GUI access to Services Director.
Custom TCP Rule	TCP	8100	Custom Remote Management IP Address.	Administrative REST API access to Services Director.
Custom TCP Rule	TCP	8000	Custom Remote Management IP Address.	Administrative GUI access to the Analytics Application

Services Director Security Group: Services Director Peer

You must add the following rules to the Security Group *used by the Services Director*.

The following rules support flows from the Services Director peer.

In these rules, the ID of the Services Director Security Group is required. This opens the specified ports to all instances that use the Security Group.

Type	Protocol	Port Range	Source	Description
Custom UDP Rule	UDP	9090	Custom Security Group ID	Internal Services Director cluster communication.
Custom TCP Rule	TCP	9070	Custom Security Group ID	Internal Services Director cluster communication (REST API).
Custom UDP Rule	UDP	9080	Custom Security Group ID	Internal Services Director cluster communication.
Custom TCP Rule	TCP	9080	Custom Security Group ID	Internal Services Director cluster communication.
Custom TCP Rule	TCP	9090	Custom Security Group ID	Internal Services Director cluster communication (GUI).
Custom UDP Rule	UDP	9091	Custom Security Group ID	Internal Services Director cluster communication.
Custom TCP Rule	TCP	8100	Custom Security Group ID	Services Director Web Service (REST API) for peer monitoring.

Type	Protocol	Port Range	Source		Description
HTTP	TCP	80	Custom	Security Group ID	Internal Services Director cluster communication.
HTTPS	TCP	443	Custom	Security Group ID	Internal Services Director cluster communication.
All ICMP - IPv4	All	N/A	Custom	Security Group ID	Ping (for monitoring).
MySQL/Aurora	TCP	3306	Custom	Security Group ID	MySQL internal (required for monitoring and failover).

Note: The following rules are also required, but you cannot add these until after you have created the Elastic IPs required the Primary and Secondary Services Director instances and the SEA, see [“Updating Security Rules for Services Director Instance IP Addresses” on page 74](#)

Type	Protocol	Port Range	Source	Description	
Custom TCP Rule	TCP	3306	Custom	EIP of the Primary Services Director instance.	MySQL inventory database access.
Custom TCP Rule	TCP	3306	Custom	EIP of the Secondary Services Director instance.	MySQL inventory database access.
Custom TCP Rule	TCP	3306	Custom	EIP of the Services Director SEA.	MySQL inventory database access.

Services Director Security Group: vTM Estate

You must add the following rules to the Security Group *used by the Services Director*.

Create the following rules for *each* vTM in the estate of the Services Director. You may be able to use a suitable IP address range that covers multiple vTMs to minimize the number of rules required.

Type	Protocol	Port Range	Source		Description
Custom TCP Rule	TCP	8100	Custom	EIP of the vTM.	vTM self-registration requests.
Custom TCP Rule	TCP	8101	Custom	EIP of the vTM.	vTM Universal FLA licensing request.
Custom TCP Rule	TCP	8102	Custom	EIP of the vTM.	Required for vTM Communications Channel, see “Working with vTM Communications Channel” on page 126.

vTM Security Group: Remote Management and Administration

You must add the following rules to the Security Group *used by individual vTMs*.

The following rules support flows from Remote Management and Administration.

In these rules, the IP address of the device from which you will perform remote management is required.

Note: For all rules, CIDR format must be used for any IP address range. This takes the form `xx.xx.xx.xx/nn`, where `nn` is the size of the subnet mask. The subnet mask size cannot be omitted, even when you specify a single IP address. That is, if you require `10.11.12.13` as an IP address, you must specify `10.11.12.13/32`.

Type	Protocol	Port Range	Source	Description
SSH	TCP	22	Custom Remote Management IP Address.	Administrative shell access to the vTM.
Custom TCP Rule	TCP	9090	Custom Remote Management IP Address.	Administrative GUI access to the vTM.
Custom TCP Rule	TCP	9070	Custom Remote Management IP Address.	Administrative REST API access to the vTM.

vTM Security Group: vTM Peers

You must add the following rules to the Security Group *used by individual vTMs*.

The following rules support flows from the vTM peers in the same cluster.

Note: This is a minimum suggested set of rules to support vTM clustering. See the *Virtual Traffic Manager documentation* for additional advice on AWS Security Groups.

In these rules, the ID of the vTM Security Group is required. This opens the specified ports to all vTMs that use the Security Group.

Type	Protocol	Port Range	Source	Description
Custom UDP Rule	UDP	9080	Custom Security Group ID	vTM internal cluster communication.
Custom UDP Rule	UDP	9090	Custom Security Group ID	vTM internal cluster communication.
Custom TCP Rule	TCP	9080	Custom Security Group ID	vTM internal cluster communication.
Custom TCP Rule	TCP	9090	Custom Security Group ID	vTM GUI.

vTM Security Group: Services Director Estate Manager

Note: These rules are not required when vTM Communications Channel is active on the vTM.

You must add the following rule to the Security Group *used by individual vTMs*.

The following rule supports flows from the Services Director Estate Manager.

In this rule, the ID of the Services Director Security Group is required.

Note: This rule is not required when vTM Communications Channel is active on the vTM.

Type	Protocol	Port Range	Source	Description
Custom TCP Rule	TCP	9070	Custom Services Director Security Group ID	vTM REST API access for configuration, backup and restore, API proxy, and so on.

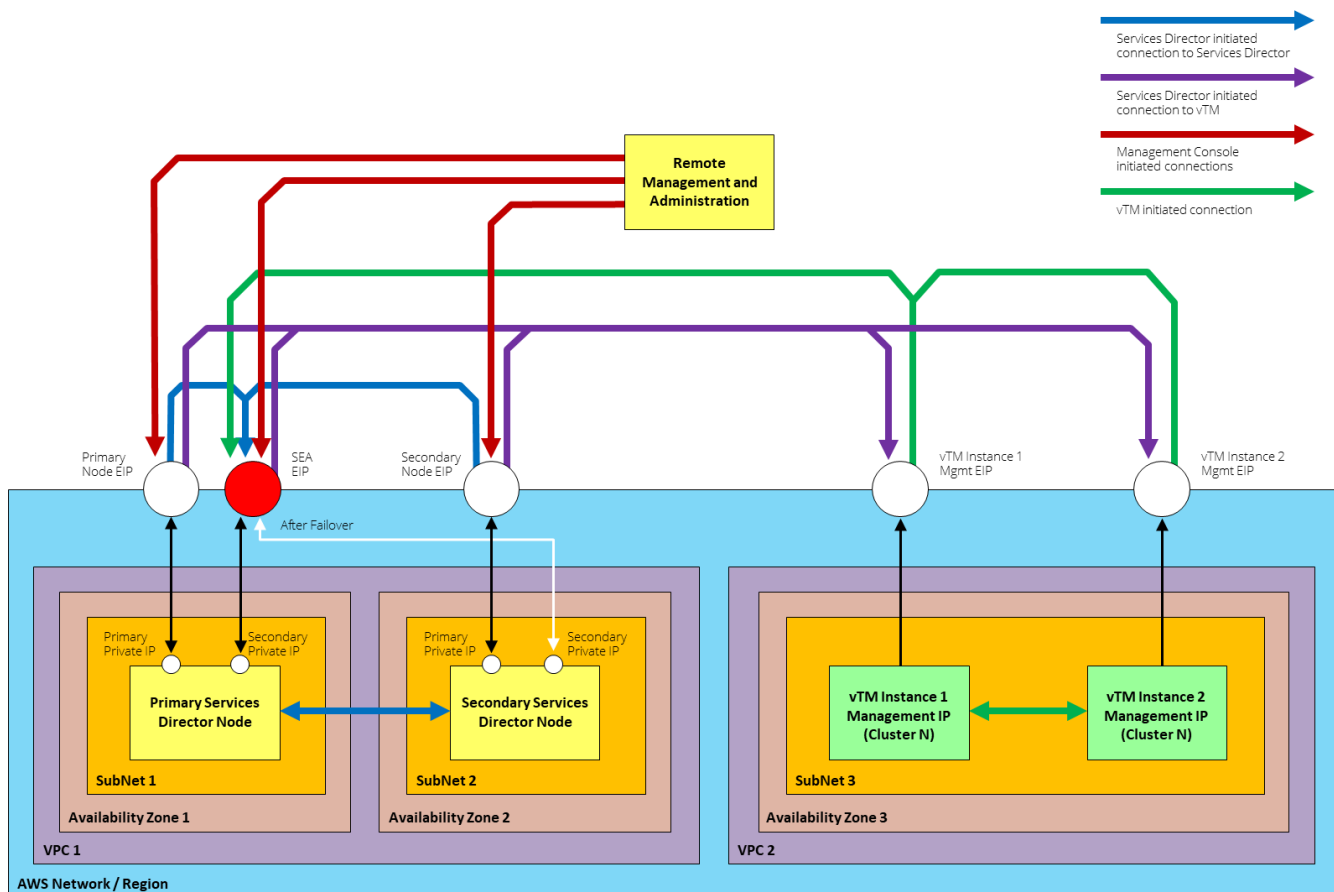
Scenario 2 - Elastic IPs with vTMs in Different VPCs

In this scenario:

- AWS Elastic IPs (EIPs) are used for the Primary and Secondary nodes, and for the Services Endpoint Address (SEA).
- The Services Director nodes are in the same AWS VPC, although they can be in separate Subnets or Availability Zones.
- The vTMs in the estate of the Services Director pair are in a different VPC (or Region).
- The self-registration feature of vTM will not work. However, manual registration of vTMs can be achieved from the Services Director by specifying a vTM's management EIP in the **Add a vTM instance** dialog, see **“Registering a Virtual Traffic Manager (Universal FLA)” on page 154**.

The flow of requests through the defined IP addresses is as follows:

FIGURE 19 Elastic IPs with vTMs in Different VPCs



Note: When vTM Communications Channel (Comms Channel) is in use, there are minor differences in the scenario diagram above. All Services Director-initiated connections to the vTM (purple in the diagram) are replaced by vTM-initiated connections to the Services Director (green in the diagram). See **“Using Scenario 2 with vTM Communications Channel” on page 48**.

Note: When you join a Secondary Services Director to a Primary Services Director, always specify the Primary Private IP Address of the Primary Services Director. This allows the majority of traffic between the Services Director nodes to be routed within the VPC and not over the public Internet.

To set up the required AWS Security Group for this scenario, you must create the rules listed in the following sections:

- **“Services Director Security Group: Remote Management and Administration” on page 49.**
- **“Services Director Security Group: Services Director Peer” on page 49.**
- **“Services Director Security Group: vTM Estate” on page 50.**
- **“vTM Security Group: Remote Management and Administration” on page 51.**
- **“vTM Security Group: vTM Peers” on page 51.**
- **“vTM Security Group: Services Director Estate Manager” on page 52.**

Note: For all rules, CIDR format must be used for each IP address range. This takes the form `xx.xx.xx.xx/nn`, where `nn` is the size of the subnet mask. The subnet mask size cannot be omitted, even when you specify a single IP address. That is, if you require `10.11.12.13` as an IP address, you must specify `10.11.12.13/32`.

Using Scenario 2 with vTM Communications Channel

When vTM Communications Channel (Comms Channel) is in use, there are minor differences from the scenario described above. All Services Director-initiated connections to the vTM (purple in the diagram) are replaced by vTM-initiated connections to the Services Director (green in the diagram).

If you are using Comms Channel, you should not attempt to connect to the Services Director from the vTM until you have its Elastic IP address. That is, you must perform the following tasks in order:

1. Create the vTM in AWS. **“Launching and Configuring the Primary Services Director on AWS” on page 28.**
2. Run the Configuration Wizard for the vTM, but do not specify Services Director details for the vTM.
3. Associate an Elastic IP address to the vTM’s Private IP address. **“Creating Elastic IP Addresses for the Services Director Instance” on page 69.**
4. Configure Security Groups on the Services Director, using the Elastic IP for the vTM (refer to the sections that follow).
5. Log into the vTM and self-register the vTM with the Services Director, see **“Requesting Self Registration on a Configured vTM” on page 177.**

Services Director Security Group: Remote Management and Administration

You must add the following rules to the Security Group *used by the Services Director*.

The following rules support flows from Remote Management and Administration.

In these rules, the IP address of the device from which you will perform remote management is required.

Note: For all rules, CIDR format must be used for any IP address range. This takes the form xx.xx.xx.xx/nn, where nn is the size of the subnet mask. The subnet mask size cannot be omitted, even when you specify a single IP address. That is, if you require 10.11.12.13 as an IP address, you must specify 10.11.12.13/32.

Type	Protocol	Port Range	Source	Description
SSH	TCP	22	Custom Remote Management IP Address.	Administrative shell access to Services Director.
HTTPS	TCP	443	Custom Remote Management IP Address.	Administrative GUI access to Services Director.
Custom TCP Rule	TCP	8100	Custom Remote Management IP Address.	Administrative REST API access to Services Director.
Custom TCP Rule	TCP	8000	Custom Remote Management IP Address.	Administrative GUI access to the Analytics Application

Services Director Security Group: Services Director Peer

You must add the following rules to the Security Group *used by the Services Director*.

The following rules support flows from the Services Director peer.

In these rules, the ID of the Services Director Security Group is required. This opens the specified ports to all instances that use the Security Group.

Type	Protocol	Port Range	Source	Description
Custom UDP Rule	UDP	9090	Custom Security Group ID	Internal Services Director cluster communication.
Custom TCP Rule	TCP	9070	Custom Security Group ID	Internal Services Director cluster communication (REST API).
Custom UDP Rule	UDP	9080	Custom Security Group ID	Internal Services Director cluster communication.
Custom TCP Rule	TCP	9080	Custom Security Group ID	Internal Services Director cluster communication.
Custom TCP Rule	TCP	9090	Custom Security Group ID	Internal Services Director cluster communication (GUI).
Custom UDP Rule	UDP	9091	Custom Security Group ID	Internal Services Director cluster communication.
Custom TCP Rule	TCP	8100	Custom Security Group ID	Services Director Web Service (REST API) for peer monitoring.

Type	Protocol	Port Range	Source		Description
HTTP	TCP	80	Custom	Security Group ID	Internal Services Director cluster communication.
HTTPS	TCP	443	Custom	Security Group ID	Internal Services Director cluster communication.
All ICMP - IPv4	All	N/A	Custom	Security Group ID	Ping (for monitoring).
MySQL/Aurora	TCP	3306	Custom	Security Group ID	MySQL internal (required for monitoring and failover).

Note: The following rules are also required, but you cannot add these until after you have created the Elastic IPs required the Primary and Secondary Services Director instances and the SEA, see [“Updating Security Rules for Services Director Instance IP Addresses” on page 74](#)

Type	Protocol	Port Range	Source		Description
Custom TCP Rule	TCP	3306	Custom	EIP of the Primary Services Director instance.	MySQL inventory database access.
Custom TCP Rule	TCP	3306	Custom	EIP of the Secondary Services Director instance.	MySQL inventory database access.
Custom TCP Rule	TCP	3306	Custom	EIP of the Services Director SEA.	MySQL inventory database access.

Services Director Security Group: vTM Estate

You must add the following rules to the Security Group *used by the Services Director*.

Create the following rules for *each* vTM in the estate of the Services Director. You may be able to use a suitable IP address range that covers multiple vTMs to minimize the number of rules required.

Note: There is no rule for self-registration in this category, as self-registration is not supported in this scenario.

Type	Protocol	Port Range	Source		Description
Custom TCP Rule	TCP	8101	Custom	EIP of the vTM.	vTM Universal FLA licensing request.
Custom TCP Rule	TCP	8102	Custom	EIP of the vTM.	Required for vTM Communications Channel, see “Working with vTM Communications Channel” on page 126 .

vTM Security Group: Remote Management and Administration

You must add the following rules to the Security Group *used by individual vTMs*.

The following rules support flows from Remote Management and Administration.

In these rules, the IP address of the device from which you will perform remote management is required.

Note: For all rules, CIDR format must be used for any IP address range. This takes the form `xx.xx.xx.xx/nn`, where `nn` is the size of the subnet mask. The subnet mask size cannot be omitted, even when you specify a single IP address. That is, if you require `10.11.12.13` as an IP address, you must specify `10.11.12.13/32`.

Type	Protocol	Port Range	Source	Description
SSH	TCP	22	Custom Remote Management IP Address.	Administrative shell access to the vTM.
Custom TCP Rule	TCP	9090	Custom Remote Management IP Address.	Administrative GUI access to the vTM.
Custom TCP Rule	TCP	9070	Custom Remote Management IP Address.	Administrative REST API access to the vTM.

vTM Security Group: vTM Peers

You must add the following rules to the Security Group *used by individual vTMs*.

The following rules support flows from the vTM peers in the same cluster.

Note: This is a minimum suggested set of rules to support vTM clustering. See the *Virtual Traffic Manager documentation* for additional advice on AWS Security Groups.

In these rules, the ID of the vTM Security Group is required. This opens the specified ports to all vTMs that use the Security Group.

Type	Protocol	Port Range	Source	Description
Custom UDP Rule	UDP	9080	Custom Security Group ID	vTM internal cluster communication.
Custom UDP Rule	UDP	9090	Custom Security Group ID	vTM internal cluster communication.
Custom TCP Rule	TCP	9080	Custom Security Group ID	vTM internal cluster communication.
Custom TCP Rule	TCP	9090	Custom Security Group ID	vTM GUI.

vTM Security Group: Services Director Estate Manager

You must add the following rule to the Security Group *used by individual vTMs*.

The following rule supports flows from the Services Director Estate Manager.

Note: You can only add these rules after you have created the Services Director AMI nodes and assigned elastic IP addresses to each and the SEA.

Note: These rules are not required when vTM Communications Channel is active on the vTM.

Type	Protocol	Port Range	Source	Description
Custom TCP Rule	TCP	9070	Custom	EIP of Primary Services Director Node. vTM REST API access for configuration, backup and restore, API proxy, and so on.
Custom TCP Rule	TCP	9070	Custom	EIP of Secondary Services Director Node. vTM REST API access for configuration, backup and restore, API proxy, and so on.
Custom TCP Rule	TCP	9070	Custom	EIP of Services Director SEA. vTM REST API access for configuration, backup and restore, API proxy, and so on.

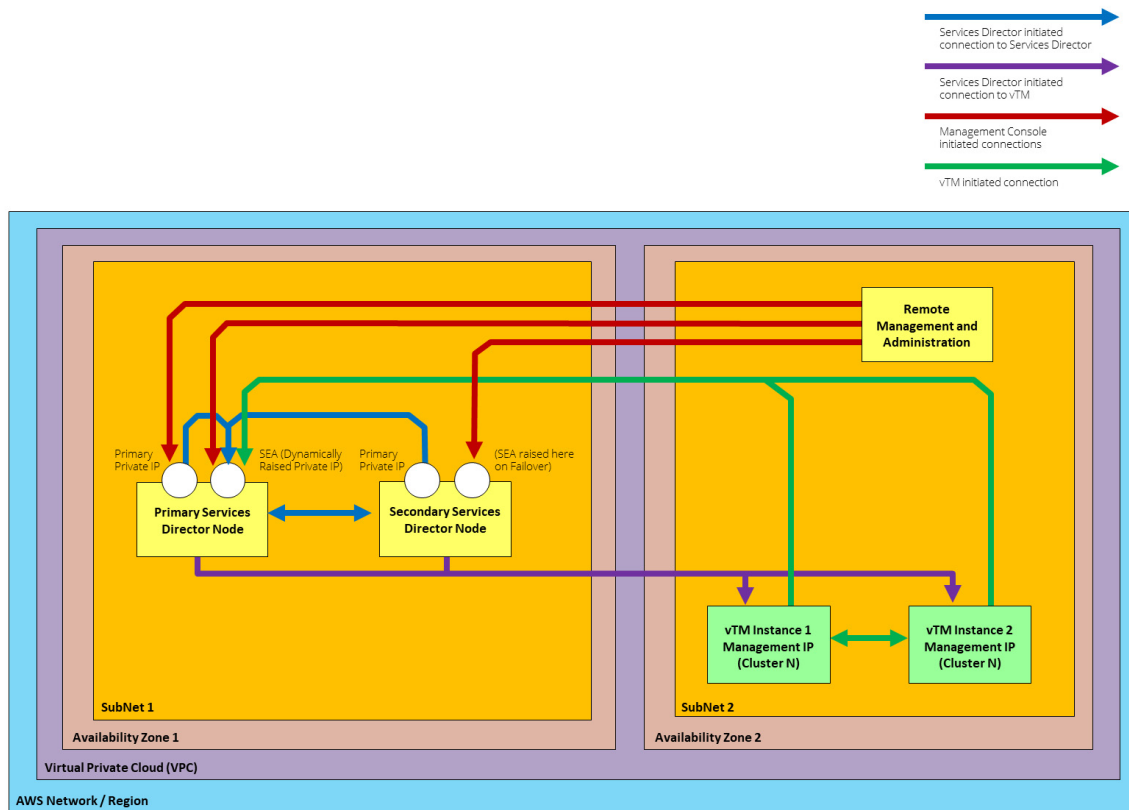
Scenario 3 - Private IPs Only

In this scenario:

- Private IP addresses are used for the Primary and Secondary nodes.
- The Services Director SEA is a private IP that can be raised on either Services Director node.
- The Primary and Secondary nodes must exist within the same Subnet as the SEA.
- The vTMs in the estate of the Services Director pair are in the same VPC, but can be in different Availability Zones or Subnets. Each uses private IP addresses only, with no elastic IP assigned for management purposes.
- All management traffic flows are directed to private IPs.
- The management console is either:
 - Inside the AWS network, within the same VPC as the Services Director nodes (as shown in the diagram), OR
 - Outside the AWS network, but able to route traffic directly to the private IP addresses within the VPC. For example, by using a peer-to-peer VPN connection from a local data centre.
- The self-registration feature is fully supported.
- The use of vTM Communications Channel is fully supported.

The flow of requests through the defined IP addresses is as follows:

FIGURE 20 Private IPs Only



Note: When vTM Communications Channel (Comms Channel) is in use, there are minor differences in the scenario diagram above. All Services Director-initiated connections to the vTM (purple in the diagram) are replaced by vTM-initiated connections to the Services Director (green in the diagram).

To set up the required AWS Security Group for this scenario, you must create the rules listed in the following sections:

- **“Services Director Security Group: Remote Management and Administration” on page 55.**
- **“Services Director Security Group: Services Director Peer” on page 55.**
- **“Services Director Security Group: vTM Estate” on page 56.**
- **“vTM Security Group: Remote Management and Administration” on page 57.**
- **“vTM Security Group: vTM Peers” on page 57.**
- **“vTM Security Group: vTM Peers” on page 57.**

Note: For all rules, CIDR format must be used for each IP address range. This takes the form `xx.xx.xx.xx/nn`, where `nn` is the size of the subnet mask. The subnet mask size cannot be omitted, even when you specify a single IP address. That is, if you require `10.11.12.13` as an IP address, you must specify `10.11.12.13/32`.

Services Director Security Group: Remote Management and Administration

You must add the following rules to the Security Group *used by the Services Director*.

The following rules support flows from Remote Management and Administration. There is no prescribed location for Remote Management.

In these rules, the IP address(es) or IP range(s) that can validly access the Services Director administration interfaces are required.

Note: For all rules, CIDR format must be used for any IP address range. This takes the form xx.xx.xx.xx/nn, where nn is the size of the subnet mask. The subnet mask size cannot be omitted, even when you specify a single IP address. That is, if you require 10.11.12.13 as an IP address, you must specify 10.11.12.13/32.

Type	Protocol	Port Range	Source	Description
SSH	TCP	22	Custom Remote Management IP Address.	Administrative shell access to Services Director.
HTTPS	TCP	443	Custom Remote Management IP Address.	Administrative GUI access to Services Director.
Custom TCP Rule	TCP	8100	Custom Remote Management IP Address.	Administrative REST API access to Services Director.
Custom TCP Rule	TCP	8000	Custom Remote Management IP Address.	Administrative GUI access to the Analytics Application

Services Director Security Group: Services Director Peer

You must add the following rules to the Security Group *used by the Services Director*.

The following rules support flows from the Services Director peer.

In these rules, the ID of the Services Director Security Group is required. This opens the specified ports to all instances that use the Security Group.

Type	Protocol	Port Range	Source	Description
Custom UDP Rule	UDP	9090	Custom Security Group ID	Internal Services Director cluster communication.
Custom TCP Rule	TCP	9070	Custom Security Group ID	Internal Services Director cluster communication (REST API).
Custom UDP Rule	UDP	9080	Custom Security Group ID	Internal Services Director cluster communication.
Custom TCP Rule	TCP	9080	Custom Security Group ID	Internal Services Director cluster communication.
Custom TCP Rule	TCP	9090	Custom Security Group ID	Internal Services Director cluster communication (GUI).
Custom UDP Rule	UDP	9091	Custom Security Group ID	Internal Services Director cluster communication.

Type	Protocol	Port Range	Source		Description
Custom TCP Rule	TCP	8100	Custom	Security Group ID	Services Director Web Service (REST API) for peer monitoring.
HTTP	TCP	80	Custom	Security Group ID	Internal Services Director cluster communication.
HTTPS	TCP	443	Custom	Security Group ID	Internal Services Director cluster communication.
All ICMP - IPv4	All	N/A	Custom	Security Group ID	Ping (for monitoring).
MySQL/Aurora	TCP	3306	Custom	Security Group ID	MySQL internal (required for monitoring and failover).

Services Director Security Group: vTM Estate

You must add the following rules to the Security Group *used by the Services Director*.

Create the following two rules for *each* vTM in the estate of the Services Director. You may be able to use a suitable IP address range that covers multiple vTMs to minimize the number of rules required.

In these rules, the ID of the vTM Security Group for is required. This opens the specified ports to all vTMs that use the Security Group.

Type	Protocol	Port Range	Source		Description
Custom TCP Rule	TCP	8100	Custom	vTM Security Group ID	vTM self-registration requests.
Custom TCP Rule	TCP	8101	Custom	vTM Security Group ID	vTM Universal FLA licensing request.
Custom TCP Rule	TCP	8102	Custom	vTM Security Group ID	Required for vTM Communications Channel, see “Working with vTM Communications Channel” on page 126.

vTM Security Group: Remote Management and Administration

You must add the following rules to the Security Group *used by individual vTMs*.

The following rules support flows from Remote Management and Administration.

In these rules, the IP address of the device from which you will perform remote management is required.

Note: For all rules, CIDR format must be used for any IP address range. This takes the form `xx.xx.xx.xx/nn`, where `nn` is the size of the subnet mask. The subnet mask size cannot be omitted, even when you specify a single IP address. That is, if you require `10.11.12.13` as an IP address, you must specify `10.11.12.13/32`.

Type	Protocol	Port Range	Source	Description
SSH	TCP	22	Custom Remote Management IP Address.	Administrative shell access to the vTM.
Custom TCP Rule	TCP	9090	Custom Remote Management IP Address.	Administrative GUI access to the vTM.
Custom TCP Rule	TCP	9070	Custom Remote Management IP Address.	Administrative REST API access to the vTM.

vTM Security Group: vTM Peers

You must add the following rules to the Security Group *used by individual vTMs*.

The following rules support flows from the vTM peers in the same cluster.

Note: This is a minimum suggested set of rules to support vTM clustering. See the *Virtual Traffic Manager documentation* for additional advice on AWS Security Groups.

In these rules, the ID of the vTM Security Group is required. This opens the specified ports to all vTMs that use the Security Group.

Type	Protocol	Port Range	Source	Description
Custom UDP Rule	UDP	9080	Custom Security Group ID	vTM internal cluster communication.
Custom UDP Rule	UDP	9090	Custom Security Group ID	vTM internal cluster communication.
Custom TCP Rule	TCP	9080	Custom Security Group ID	vTM internal cluster communication.
Custom TCP Rule	TCP	9090	Custom Security Group ID	vTM GUI.

vTM Security Group: Services Director Estate Manager

Note: These rules are not required when vTM Communications Channel is active on the vTM.

You must add the following rule to the Security Group *used by individual vTMs*.

The following rule supports flows from the Services Director Estate Manager.

In this rule, the ID of the Services Director Security Group is required.

Type	Protocol	Port Range	Source	Description
Custom TCP Rule	TCP	9070	Custom Services Director Security Group ID	vTM REST API access for configuration, backup and restore, API proxy, and so on. Note: This rule is not required when vTM Communications Channel is active on the vTM.

Creating an AWS Security Group

Once you have gathered all required information for your Services Director scenario, you can create the required AWS Security Group.

To create an AWS Security Group:

1. Login to the *AWS Management Console*.
2. On the top bar of the AWS Management Console, select the required **Region**.
3. On the AWS top bar, click **Services** and then locate the **Network & Content Delivery** options.
4. Under **Network & Content Delivery**, select **VPC**.

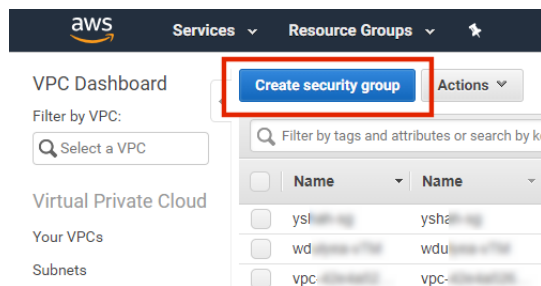
The **VPC Dashboard** appears.

5. In the left menu, Under **Security**, click **Security Groups**.

A list of your existing Security Groups appears.

6. Examine your Security Groups and decide if you have one that matches your networking requirements for your Services Director nodes. If there is a Security Group, no further actions is required, and this process is complete.
7. Click **Create security group**.

FIGURE 21 Create Security Group



The **Create security group** page appears.

FIGURE 22 Create Security Group

[Security Groups](#) > Create security group

Create security group

A security group acts as a virtual firewall for your instance to control inbound and outbound traffic. To create a new security group fill in the fields below.

Security group name* ⓘ

Description* ⓘ

VPC ⓘ

* Required

[Cancel](#) [Create](#)

- Specify your required details and click **Create**.

A confirmation message appears. For example:

FIGURE 23 Create Security Group: Success

[Security Groups](#) > Create security group

Create security group

✓ The following security group was created:

Security Group ID [sg-0b4ba97b04c2649fc](#)

[Close](#)

Click the Security Group ID link to show the new Security Group in the **Security Groups** list. There are no inbound rules defined on this new Security Group. For example:

FIGURE 24 New Security Group

VPC Dashboard

Filter by VPC:

Virtual Private Cloud

Your VPCs

Subnets

Route Tables

Internet Gateways

Egress Only Internet Gateways

DHCP Options Sets

Elastic IPs

[Create security group](#) [Actions](#)

Group ID : sg-0b4ba97b04c2649fc Add filter

Name	Group ID	Group Name	VPC ID	Type	Description
	sg-0b4ba97b04c2649fc	jk-sec-01	vpc-ec1da988	EC2-VPC	Personal security group

Security Group: [sg-0b4ba97b04c2649fc](#)

Description Inbound Rules Outbound Rules Tags

Group ID sg-0b4ba97b04c2649fc

VPC ID vpc-ec1da988

Owner 8151814

Outbound rule count 1

Group Name jk-sec-01

Description Personal security group

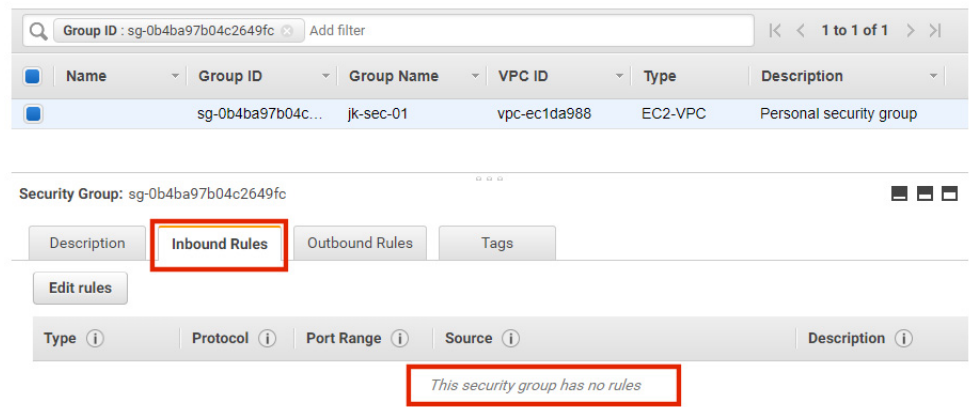
Inbound rule count 0

- Record the **Group ID** for the Security Group. This is required when adding rules to the group.

10. Click the **Inbound Rules** tab.

An empty list of inbound rules appears. For example:

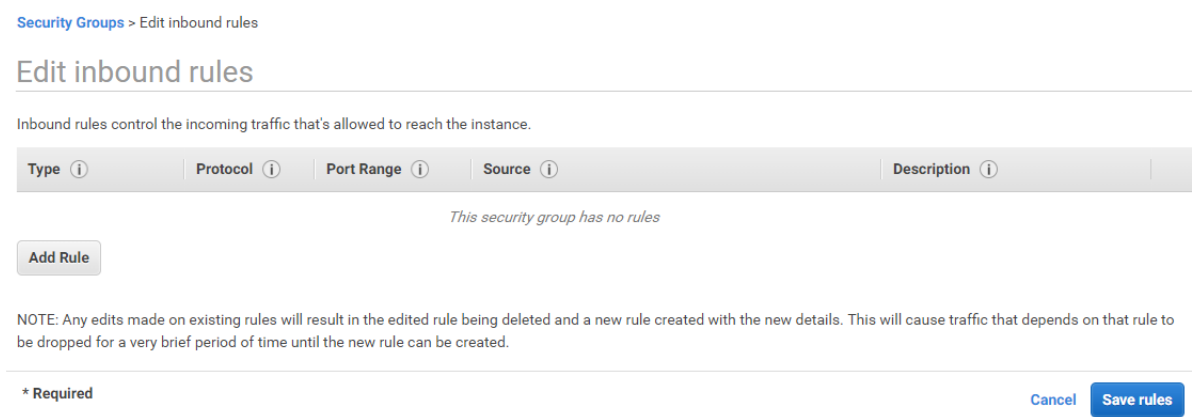
FIGURE 25 Empty Inbound Rules



11. Click **Edit Rules**.

The **Edit inbound rules** page appears. For example:


FIGURE 26 Edit Inbound Rules



12. Click **Add Rule**.

A new entry is added to the list of rules. For example:

FIGURE 27 Inbound Rules Entry



The required rules for your Security Group are described in one of the following three scenarios:

- **“Scenario 1 - Elastic IPs with vTMs in the Same VPC” on page 41.**
- **“Scenario 2 - Elastic IPs with vTMs in Different VPCs” on page 47.**
- **“Scenario 3 - Private IPs Only” on page 53.**

13. In the new entry, specify a required inbound rule and click **Add Rule**.

The new rule is added.

14. Repeat step 13 until you have recorded all required rules.

For example:

FIGURE 28 Populated Services Director Rules

Edit inbound rules

Inbound rules control the incoming traffic that's allowed to reach the instance.

Type ⓘ	Protocol ⓘ	Port Range ⓘ	Source ⓘ	Description ⓘ	
HTTP	TCP	80	Custom sg-0411f77c014	HTTP	ⓧ
Custom UDP Rule	UDP	9090	Custom sg-0411f77c014	Internal vTM communications	ⓧ
Custom TCP Rule	TCP	9070	Custom sg-0411f77c014	vTM REST API	ⓧ
Custom TCP Rule	TCP	9090	Custom 212.	Internal vTM GUI	ⓧ
Custom TCP Rule	TCP	9090	Custom sg-0411f77c014	Internal vTM communications	ⓧ
HTTPS	TCP	443	Custom 212.	HTTPS	ⓧ
HTTPS	TCP	443	Custom sg-0411f77c014	HTTPS	ⓧ
All ICMP - IPv4	ICMP	All	Custom sg-0411f77c014	Internal SD monitoring	ⓧ
Custom TCP Rule	TCP	8101	Custom sg-0411f77c014	vTMs Universal FLA	ⓧ
Custom TCP Rule	TCP	9080	Custom sg-0411f77c014	Internal vTM communications	ⓧ
SSH	TCP	22	Custom 212.	SSH (for CLI)	ⓧ
Custom TCP Rule	TCP	8100	Custom 212.	SD REST API	ⓧ
Custom TCP Rule	TCP	8100	Custom sg-0411f77c014	Internal SD monitoring	ⓧ
Custom UDP Rule	UDP	9080	Custom sg-0411f77c014	Internal vTM communications	ⓧ
MYSQL/Aurora	TCP	3306	Custom sg-0411f77c014	Internal MySQL communications	ⓧ
Custom UDP Rule	UDP	9091	Custom sg-0411f77c014	Internal vTM communications	ⓧ

Add Rule

NOTE: Any edits made on existing rules will result in the edited rule being deleted and a new rule created with the new details. This will cause traffic that depends on that rule to be dropped for a very brief period of time until the new rule can be created.

* Required

Cancel Save rules

15. Once you have added all rules, click **Save Rules**.

A confirmation message appears.

16. Click **Close**.

The rules are shown in the **Inbound Rules** tab.

Once you have a Security Group with all required rules, you can launch a Services Director instance, see [“Launching a Services Director AMI Instance on AWS” on page 64](#).

Note: You will need to add additional rules to the Security Group for any elastic IP addresses requested/assigned to each Services Director node and to the SEA, see [“Updating Security Rules for Services Director Instance IP Addresses” on page 74](#).

Launching a Services Director AMI Instance on AWS

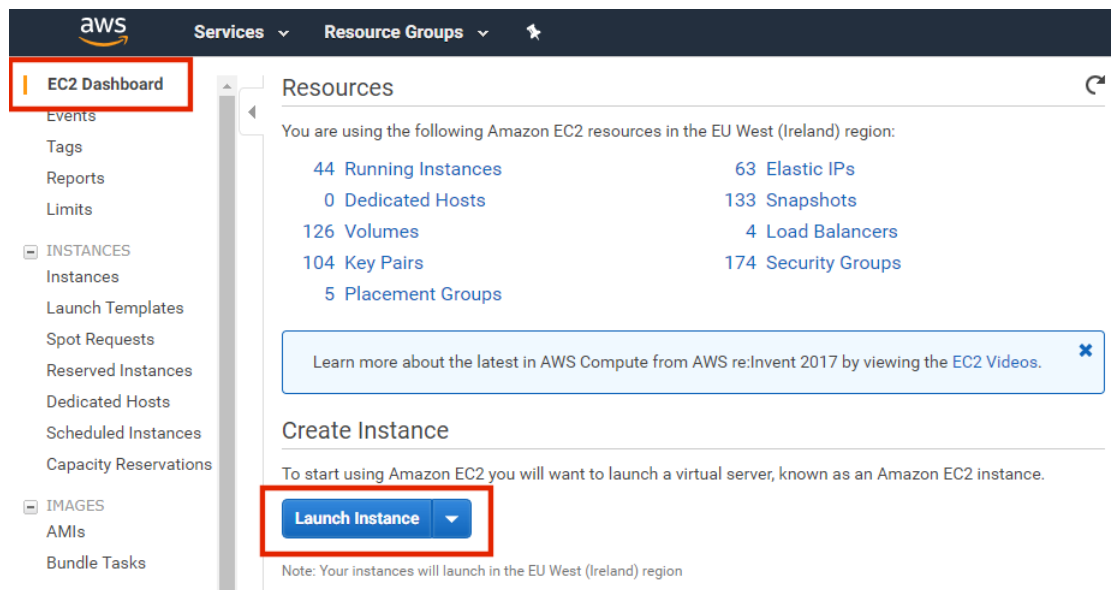
Once you have made all required preparations, you can launch a Services Director AMI instance on AWS. The processes for Primary and Secondary Services Directors are the same.

To launch a Services Director AMI instance:

1. Login to the *AWS Management Console*.
2. On the top bar of the AWS Management Console, select the required **Region**.
3. On the AWS top bar, click **Services** and then locate the **Compute** options.
4. Under **Compute**, select **EC2**.

The **EC2 Dashboard** appears. For example:

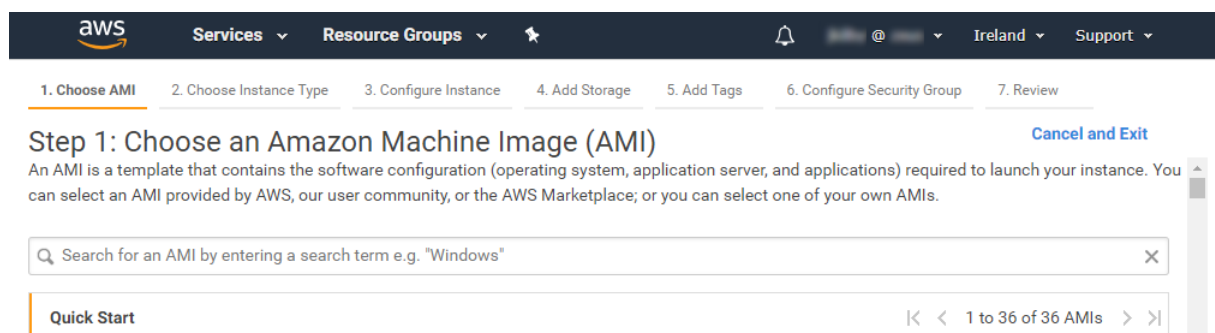
FIGURE 29 EC2 Dashboard



5. Under **Create Instance**, click **Launch Instance**.

The first panel of the AMI Launch Wizard appears. For example:

FIGURE 30 Launch Wizard 1: Select AMI



- In the Launch Wizard, locate the Services Director AMI and click **Select**.

The next panel of the AMI Launch Wizard (**Choose Instance Type**) appears.

- On the **Choose Instance Type** panel, select a General Purpose *T2.large* Services Director AMI, or a better specification.
- Click **Configure Instance Details**.

The next panel of the AMI Launch Wizard (**Configure Instance**) appears.

- On the **Configure Instance** panel, set the following properties:
 - Number of Instance:** Select *1*.
 - Network:** Select the required AWS VPC for your Services Director. You prepared this earlier, see [“Creating an AWS Virtual Private Cloud” on page 35](#).
 - Subnet:** Select the required AWS Subnet from your selected VPC. You prepared this earlier, see [“Creating AWS Subnets” on page 37](#).
 - Auto-assign Public IP:** Select *Enable*.
 - IAM Role:** (Optional) Select your IAM role. If you do not specify **IAM Role**, you will require AWS credentials when you run the Services Director VA Setup Wizard.

For example:

FIGURE 31 Launch Wizard 2: Choose Instance Type

The screenshot displays the 'Step 3: Configure Instance Details' panel of the AWS IAM Launch Wizard. At the top, a progress bar indicates the current step is '3. Configure Instance', with previous steps '1. Choose AMI', '2. Choose Instance Type', and '4. Add Storage' also visible. Below the progress bar, the title 'Step 3: Configure Instance Details' is followed by a descriptive text: 'Configure the instance to suit your requirements. You can launch multiple instances from the same AMI, request Spot instances, assign an IAM management role to the instance, and more.'

The configuration options are as follows:

- Number of instances:** A text input field containing the value '1'. To its right is a link 'Launch into Auto Scaling Group' with an information icon.
- Purchasing option:** A section with a checkbox labeled 'Request Spot instances', which is currently unchecked.
- Network:** A dropdown menu showing 'vpc-989 | SD-A'. To its right is a link 'Create new V' with a plus icon.
- Subnet:** A dropdown menu showing 'subnet-26a3 | SD-A | eu-west-1a'. Below this, it states '65518 IP Addresses available'. To its right is a link 'Create new s' with a plus icon.
- Auto-assign Public IP:** A dropdown menu showing 'Enable'.
- Placement group:** A section with a checkbox labeled 'Add instance to placement group.', which is unchecked.
- Capacity Reservation:** A dropdown menu showing 'Open'. To its right is a link 'Create new C' with a plus icon.
- IAM role:** A dropdown menu showing 'None'. To its right is a link 'Create new IA' with a plus icon.

A red rectangular box highlights the 'Network', 'Subnet', 'Auto-assign Public IP', and 'IAM role' sections, indicating the specific configuration choices made for the instance.

- Click **Add Storage**.

The next panel of the AMI Launch Wizard (**Add Storage**) appears.

Note: By default, there are no required changes on the **Add Storage** panel.

11. On the **Add Storage** panel, change the storage options as required.

12. Click **Add Tags**.

The next panel of the AMI Launch Wizard (**Add Tags**) appears.

13. (Optional) On the **Add Tags** panel, create any tags that are required.

14. Click **Configure Security Group**.

The next panel of the AMI Launch Wizard (**Configure Security Group**) appears.

15. On the **Configure Security Group** panel, for **Assign a security group**, select the *Select an existing security group* option.

A list of available AWS Security Groups appears.

16. Select the Security Group that you prepared for your Services Director nodes, see [“Preparing an AWS Security Group” on page 40](#).

17. Click **Review and Launch**.

The final panel of the AMI Launch Wizard (**Review Instance Launch**) appears.

18. On the **Review Instance Launch** panel, confirm all details for your AMI instance and (optionally) go back through the wizard to make any final changes.

19. Click **Launch**.

The **Select an existing key pair or create a new key pair** dialog appears. For example:

FIGURE 32 Launch Wizard: Key Pair

Select an existing key pair or create a new key pair X

A key pair consists of a **public key** that AWS stores, and a **private key file** that you store. Together, they allow you to connect to your instance securely. For Windows AMIs, the private key file is required to obtain the password used to log into your instance. For Linux AMIs, the private key file allows you to securely SSH into your instance.

Note: The selected key pair will be added to the set of keys authorized for this instance. Learn more about [removing existing key pairs from a public AMI](#).

Choose an existing key pair ▼

Select a key pair

admin ▼

☐ I acknowledge that I have access to the selected private key file (admin.pem), and that without this file, I won't be able to log into my instance.

Cancel Launch Instances

20. In this dialog, either:

- Select an existing key pair for which you have the private key.
- Create a new key pair. Once you have created the pair, you will need to download the private key.


The public key will be embedded in the Services Director instance.

The private key must be retained for reference. It is required to retrieve the default password for the Services Director instance, see [“Retrieving the Default Password from the Services Director Instance Using SSH” on page 75](#).

21. Click **Launch Instances**.

The Services Director instance launches, and a confirmation appears. For example:


FIGURE 33 Launch Status



22. Click the instance ID link.

23. The new Services Director instance is listed on the **Instances** panel of the **EC2 Dashboard**.


FIGURE 34 New AMI Services Director



You do not have to wait for the instance to complete its initialization.

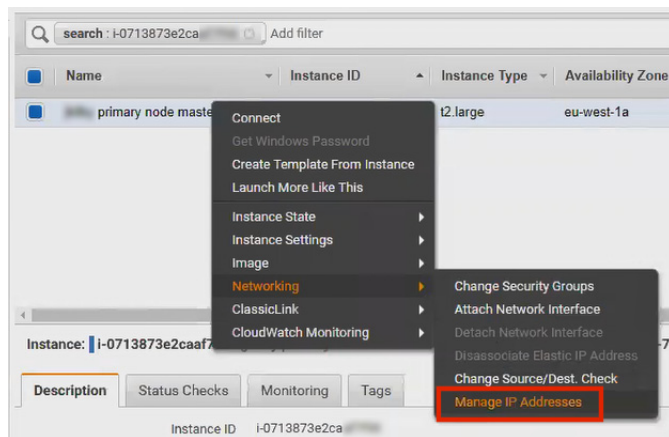
24. (Optional) Add a **Name** for the Services Director instance by hovering the mouse pointer over the empty **Name** property and clicking the **Edit** icon. For example:

FIGURE 35 Editing the Services Director Instance Name



25. Right click on the Services Director instance, and select **Networking** and then click **Manage IP Addresses**.

FIGURE 36 Services Director Instance Networking Menu

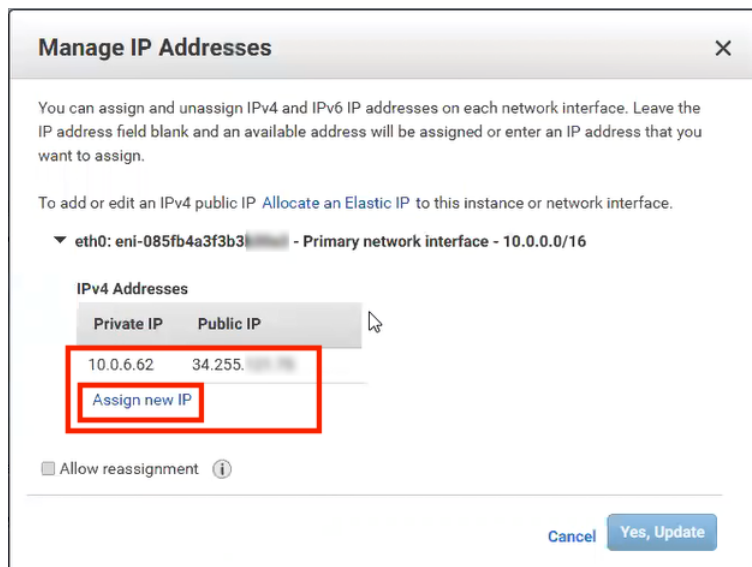


The **Manage IP Addresses** dialog appears. This shows:

- The Primary **Private IP** assigned to the Services Director instance.
- The **Public IP** (non-elastic) address requested during Step 3 of the AMI Launch Wizard.

For example:

FIGURE 37 Manage IP Addresses



Once you have launched the Services Director instance, you can optionally configure elastic IP addresses for the Services Director node and its SEA, see **“Creating Elastic IP Addresses for the Services Director Instance” on page 69**.

If you intend to use Private IP Addresses only (for example, to access your Services Directors via a VPN tunnel from your local network to the AWS Network), you can continue from [“Retrieving the Default Password for a Services Director Instance” on page 74](#).

Creating Elastic IP Addresses for the Services Director Instance

This section describes an optional process to add elastic IP addresses to the Services Director instance. This process is required if you need Internet-facing IP addresses for the Services Director node and its SEA.

Note: The processes for Primary and Secondary Services Director instances are very similar. You do not create the elastic SEA on the Secondary instance (see below).

If you intend to use Private IP Addresses only for example, to access your Services Directors via a VPN tunnel from your local network to the AWS Network), you can continue from [“Retrieving the Default Password for a Services Director Instance” on page 74](#).

You must create either one or two elastic IP addresses:

- The *first* elastic IP address will be the Internet-facing IP address of the Services Director node. You must associate this with the Primary Private IP address of the node.
- The *second* elastic IP address will be the SEA for the Services Director pair.

Note: You only need to create the elastic IP address for the SEA from the Primary instance, as a single SEA will be shared by your Primary and Secondary nodes. This is the only difference between the processes for the Primary and Secondary Services Director instances.

- To use an elastic SEA, you must also prepare a Secondary Private IP address on both Primary and Secondary instances, but you *do not* associate this to the elastic SEA. In operation, the elastic SEA always directs requests to the Secondary Private IP of the *Active* node automatically.

To configure this, perform the following steps:

1. Login to the *AWS Management Console*.
2. On the top bar of the AWS Management Console, select the required **Region**.
3. On the AWS top bar, click **Services** and then locate the **Compute** options.
4. Under **Compute**, select **EC2**. The **EC2 Dashboard** appears.
5. In the left menu, under **Instances**, select **Instances**. A list of your instances appears.
6. Locate your Services Director instance.

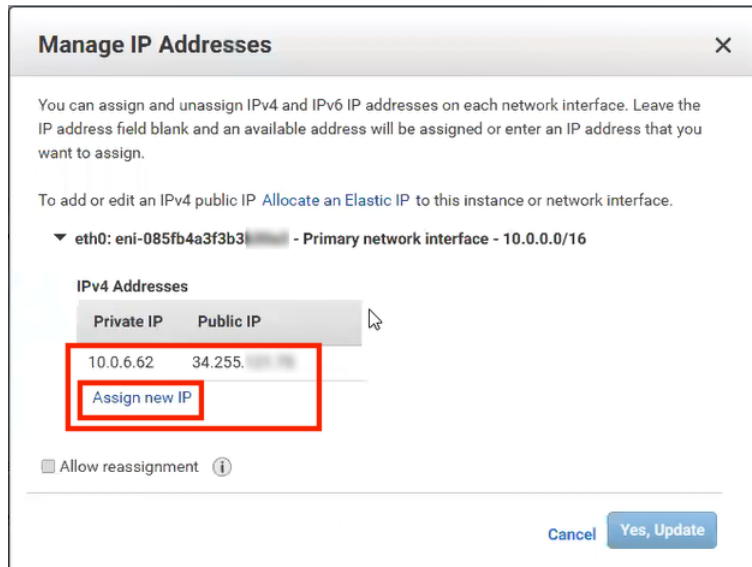
- Right click on the Services Director instance, and select **Networking** and then click **Manage IP Addresses**.

The **Manage IP Addresses** dialog appears. This shows:

- The Primary **Private IP** assigned to the Services Director instance.
- The **Public IP** (non-elastic) address requested during Step 3 of the AMI Launch Wizard.

For example:

FIGURE 38 Manage IP Addresses



- Click **Assign new IP** below the **Private IP** to create a Secondary Private IP address.
- Click **Yes, Update**.

The Secondary Private IP Address appears. For example:

FIGURE 39 Manage IP Addresses



- Click **Allocate an Elastic IP**.

The **Allocate new address** page appears.


11. Select a **Scope** of *VPC* and click **Allocate**.

FIGURE 40 Allocate New Address

Addresses > Allocate new address

Allocate new address

Allocate a new Elastic IP address by selecting the scope in which it will be used


Scope ☒ VPC ☐ Classic 

[Cancel](#) [Allocate](#)

A confirmation message of the new elastic IP address appears.

FIGURE 41 Allocate New Address Confirmation

Allocate new address

 **New address request succeeded**

Elastic IP 63.33.

[Close](#)

12. Click the **Elastic IP** link, and (optionally) on the list of elastic IPs, add a name to the new elastic IP.

Note: The elastic IP now exists, but is not yet associated with the Services Director instance.

13. Right click on the elastic IP, and click **Associate Address**.

FIGURE 42 Associate Address

Elastic IP: 63.33. Add filter

Name	Elastic IP	Allocation ID	Instance	Private IP address
primary	63.33.			

Release addresses
Associate address
 Disassociate address
 Move to VPC scope
 Restore to EC2 scope
 Add/Edit Tags

Note: The **Instance** and **Public IP address** properties for the elastic IP are not yet set.

14. On the **Associate Address** page:

- Select the Services Director **Instance** you have just created.
- Select the Primary **Private IP** address for the selected Services Director instance.

FIGURE 43 Associate Address

Associate address

Select the instance OR network interface to which you want to associate this Elastic IP address (63.33. [redacted])

Resource type ☒ Instance ? ☐ Network interface

Instance ↻

Private IP ↻ ?

Reassociation ☐ Allow Elastic IP to be reassociated if already attached ?

Warning

If you associate an Elastic IP address with your instance, your current public IP address is released. [Learn more.](#)

[Cancel](#) [Associate](#)

15. Click **Associate**.

A confirmation message appears.

16. Click **Close** and return to the list of elastic IPs.

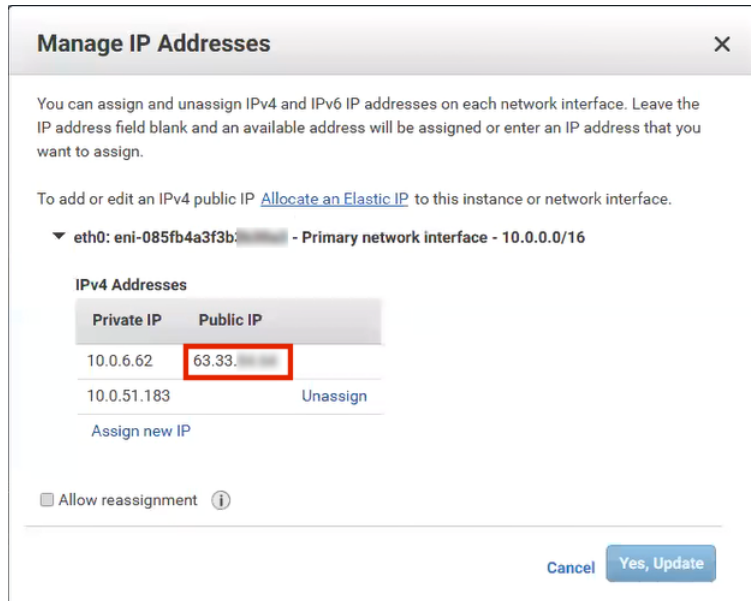
FIGURE 44 Manage IP Addresses Complete

Elastic IP : 63.33. [redacted] + Add filter						
<input checked="" type="checkbox"/>	Name	Name	Elastic IP	Allocation ID	Instance	Private IP address
<input checked="" type="checkbox"/>	primary	primary	63.33. [redacted]	eipalloc-0afa1[redacted]	i-0713873e2ca[redacted]	10.0.6.62

17. Return to the **Manage IP Addresses** dialog.

This dialog now contains all required IP addresses, including the elastic IP address associated with the Primary Private IP, which is listed as the **Public IP**. For example:

FIGURE 45 Manage IP Addresses: First Elastic



18. On the Primary Services Director instance only, click **Allocate an Elastic IP** again to create a second elastic IP. This will be used as the SEA when you use the Setup Wizard to install and configure the Primary Services Director node.

Note: Do *not* associate this second elastic IP with the Secondary Private IP address. Services Director will perform this automatically to always direct requests to the *Active* node.

19. Close the dialog to conclude the creation and association of elastic IP addresses for the Services Director instance.

Once you have launched the Services Director instance and configured its IP addresses, you must update your Security Group to add rules for these IP addresses, see **“Updating Security Rules for Services Director Instance IP Addresses” on page 74**.

Updating Security Rules for Services Director Instance IP Addresses

Once you have configured the IP addresses on the Primary Services Director instance, you must add these IP addresses as rules in the Security Group assigned to the Services Director.

Note: For all rules, CIDR format must be used for each IP address range. This takes the form `xx.xx.xx.xx/nn`, where `nn` is the size of the subnet mask. The subnet mask size cannot be omitted, even when you specify a single IP address. That is, if you require `10.11.12.13` as an IP address, you must specify `10.11.12.13/32`.

The additional rules you must add are shown below:

Type	Protocol	Port Range	Source	Description
MySQL/Aurora	TCP	3306	Custom	The Public IP address (typically an elastic IP) for the Services Director.
MySQL/Aurora	TCP	3306	Custom	The intended SEA (typically an elastic IP) for the Services Director instance.

The general process for adding rules is described in [“Preparing an AWS Security Group” on page 40](#).

Once you have updated the security rules to include the IP addresses of the Services Director instance, you must retrieve its default password, see [“Retrieving the Default Password for a Services Director Instance” on page 74](#).

Note: You will also need to add the Services Director IP addresses as rules in the Security Group used for each vTM in the estate of your Services Director pair.

Retrieving the Default Password for a Services Director Instance

Before you can access a Services Director instance and run the Setup Wizard, you must retrieve the password for the instance.

You can do this in two ways:

- On the AWS management console, examine the startup logs for the Services Director instance. The default password is recorded in this log, see [“Retrieving the Default Password from AWS Startup Logs” on page 75](#).
- On another machine, SSH into the Services Director instance, using the default user and the private key that you have stored on the machine. From there, the password can be retrieved using the CLI for the machine, see [“Retrieving the Default Password from the Services Director Instance Using SSH” on page 75](#).

After you have recorded the default password, you can access the Services Director instance for the first time from your browser, see [“Accessing your Services Director Instance for the First Time” on page 77](#).

Retrieving the Default Password from AWS Startup Logs

You can retrieve the default password for a Services Director instance by examining its startup logs in AWS. The password will be present in the system log after the Services Director instance fully initializes.

To do this:

1. Login to the *AWS Management Console*.
2. On the top bar of the AWS Management Console, select the required **Region**.
3. On the AWS top bar, click **Services** and then locate the **Compute** options.
4. Under **Compute**, select **EC2**. The **EC2 Dashboard** appears.
5. In the left menu, under **Instances**, select **Instances**. A list of your instances appears.
6. Locate your Services Director instance.
7. Right click on the Services Director instance, and select **Instance Settings** and then click **Get System Log**.
8. Search the system log until you locate the following section:

```
-----
Pulse Secure Services Director, version 18.3.0
-----
```

```
Welcome to Pulse Secure Services Director.
The appliance has now booted. To manage, please use a web browser
to access this URL:
```

```
Administration interface: https://xx.xx.xx.xx
Username: admin
Temporary Password: 4PGTdldzn9AwZn7
```

9. Record the Temporary Password value. This is the required default password.
10. Close the System Log.

After you have recorded the default password, you can access the Services Director instance for the first time from your browser, see [“Accessing your Services Director Instance for the First Time” on page 77](#).

Retrieving the Default Password from the Services Director Instance Using SSH

You can retrieve the default password from a Services Director instance directly. This requires the use of SSH and the private key for the Service Director instance.

Note: The public key for the selected key pair was embedded in the Services Director instance during the launch of the instance, see [“Launching a Services Director AML Instance on AWS” on page 64](#).

Note: Ensure that the permissions on your private key file conform to the instructions on the AWS Management Console.

To retrieve the default password from the Services Director instance:

1. Log into the machine where you have the private key stored.
2. Using your preferred SSH tool, SSH into the Services Director instance.

```
ssh -i <path/> admin@<ip_address>
```

In this example:

- *path* is the relative path from the current directory to the directory containing the private key file.
- *file* is the name of the private key file, which typically uses a *.pem* suffix.
- *ip_address* is the IP address for the Services Director instance. Typically, this is the elastic IP address associated with the Services Director instance.

You are then logged into the *admin* user on the Services Director instance.

```
Pulse Secure Services Director
Last Login: <timestamp>
Pulse Secure Services Director configuration wizard
Do you want to use the wizard for initial configuration?
```

3. Either:

- Respond *no* to bypass the configuration wizard and go straight to the command line.
- Respond *yes* to run the initial configuration. For AWS, this enables you to set the hostname for the instance. For example:

```
Do you want to use the wizard for initial configuration? yes
```

```
Step 1. Hostname? [current-hostname] <new-hostname>
```

```
You have entered the following information:
```

```
1. Hostname: <example-hostname>
```

```
To change an answer, enter the step number to return to.
```

```
Otherwise hit <enter> to save changes and exit.
```

```
To continue setup, navigate your web browser to the address configured above
```

```
Choice: <enter>
```

```
Configuration changes saved.
```

```
To return to the wizard from the CLI, use the "configuration jump-start"
command in configure mode. Enter configuration mode using commands "enable"
and "configure terminal". Launching CLI
```

```
<example-hostname> >
```

4. From the `<hostname>` command prompt, start configuration mode:

```
<hostname> > enable
<hostname> # configure terminal
<hostname> (config) #
```

5. Run the following CLI command:

```
<hostname> (config) # support show default-password
```

Note: This command is not listed in the command directory, and must be typed in full.

The password is then displayed.

6. Record the default password.
7. Close the SSH session.

After you have recorded the default password, you can access the Services Director instance for the first time from your browser, see [“Accessing your Services Director Instance for the First Time” on page 77](#).

Accessing your Services Director Instance for the First Time

Once you have retrieved the default password for a Services Director instance, you can log into the instance for the first time.

To access your Services Director instance:

1. In a browser window, access the IP address for the Services Director instance.

Typically, this will be the elastic IP address assigned to the instance.

Note: Do *not* use your intended SEA, as this is not associated with the instance at this point.

2. Accept the End User License Agreement (EULA).

The Services Director login page appears.

3. Log into the Services Director.

The default administration user name is *admin*, and the password is the default password you retrieved earlier, see [“Retrieving the Default Password for a Services Director Instance” on page 74](#).

Once you are logged in, the Services Director Setup Wizard starts automatically, see [“Running the Services Director VA Setup Wizard” on page 79](#).

Once you have completed the Setup Wizard, the creation of the Services Director node is complete.

Launching and Configuring the Secondary Services Director on AWS

Perform the following procedure to launch and configure a Secondary Services Director VA on AWS:

Note: The preparatory stages that were required for the Primary Services Director do not need to be repeated.

1. Launch a Services Director instance on AWS from the Services Director AMI, see [“Launching a Services Director AMI Instance on AWS” on page 64](#).
2. (Optional) Add and configure elastic IP addresses for the Secondary Services Director instance, see [“Creating Elastic IP Addresses for the Services Director Instance” on page 69](#).
3. Update your AWS Security Group to include all allocated IP addresses for the Secondary Services Director instance, see [“Updating Security Rules for Services Director Instance IP Addresses” on page 74](#).
4. Retrieve the password for the Secondary Services Director from AWS, see [“Retrieving the Default Password for a Services Director Instance” on page 74](#).
5. Access the Secondary Services Director instance using a browser, see [“Accessing your Services Director Instance for the First Time” on page 77](#).
6. Use the Setup Wizard (which starts automatically) to create your Secondary Services Director node. During this process you will join the Secondary node to the existing Primary node, see [“Running the Services Director VA Setup Wizard” on page 79](#).

Once this process is complete, your Services Director HA pair is complete.

Running the Services Director VA Setup Wizard

- **Overview: Setup Wizard** 79
- **Installing and Configuring a Primary Services Director** 80
- **Installing and Configuring a Secondary Services Director** 105
- **Accessing your Services Director VA** 110

Overview: Setup Wizard

After you have created/launched a Services Director VA on the required platform, you configure the Services Director VA using the Setup Wizard. The Setup Wizard enables you to:

- Select the role for this Services Director. That is, either *Primary* or *Secondary*.
 - A Primary Services Director can run as a standalone node, and assumes an active role in managing services.
 - A Secondary Services Director is joined to the Primary Services Director and can be promoted to the active role in the event of a failure.

When a Secondary Services Director is joined to the Primary Services Director in the Setup Wizard, a High Availability (HA) pair is formed.

- Specify a Service Endpoint Address for the Services Director.

Note: If the Service Endpoint Address is in a private network behind a NAT device, you must specify both the internal and external IP addresses for the Service Endpoint Address.
- Select whether to manage your Services Director (and vTM instances) using DNS hostnames or IP addresses. The option you choose depends on your deployment environment.
- Establish your licenses. This includes the Services Director License, plus any additional Resource Licenses (for bandwidth and analytics). These are required to complete the setup of the Services Director.
- (AWS only) Specify your AWS credentials, where an IAM role has not been defined.
- Define a master password. This password is used to encrypt the administration passwords of all Virtual Traffic Managers (vTMs).

The Setup Wizard automatically starts the first time you log in to the Services Director VA with a browser.

Note: The Setup Wizard is also used during recovery after a Services Director failure. For details, refer to the *Pulse Services Director Advanced User Guide*.

Installing and Configuring a Primary Services Director

To install and configure a Primary Services Director, perform the following procedure:

1. Start the Setup Wizard process, see [“Starting the Setup Wizard” on page 80](#).
2. Define a Service Endpoint Address (SEA), see [“Defining a Service Endpoint Address” on page 89](#).
3. Redeem a license token, see [“Redeeming a License Token” on page 92](#).
4. Generate a self-signed SSL certificate, see [“Generating a Self-Signed SSL Certificate” on page 93](#).
5. Add certificates and licenses, see [“Adding Certificates and Licenses” on page 95](#).
6. Complete the installation, see [“Completing the Services Director Installation” on page 103](#).

Starting the Setup Wizard

When you log into your Services Director for the first time, the Services Director VA Setup Wizard starts automatically.

1. Access your Services Director VA in a browser window using its IP address. Typically, this will be the elastic IP address assigned to the node.

Note: Do *not* use your intended SEA, as this is not associated with the instance at this point.

An End User License Agreement (EULA) statement appears.

FIGURE 46 Setup Wizard: EULA Page

Use of this Pulse Secure product is subject to the Pulse Secure
End User License Agreement available at
<https://www.pulsesecure.net/support/eula>.

Clicking on the "I agree" button below constitutes your acceptance
of these terms.

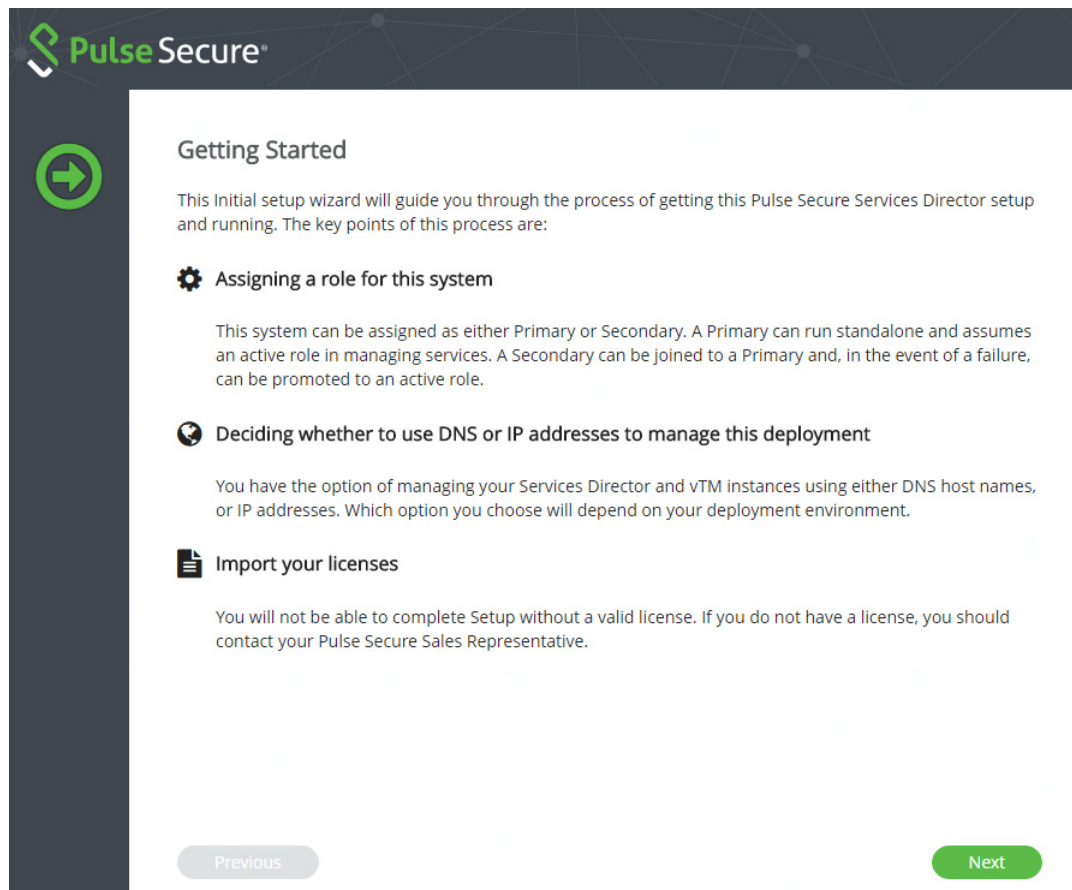


2. Click **I agree** to continue.
3. Log in using the default admin user (*admin*) and the default password.
 - For vSphere and KVM, the default password is *password*.
 - For AWS, the default password is the one you retrieved from the Services Director instance, see [“Retrieving the Default Password for a Services Director Instance” on page 74](#).

4. Click **Sign In**.

The Setup Wizard starts automatically.

FIGURE 47 Setup Wizard: Getting Started Page



5. Click **Next**.

The **Set Administration Credentials** page appears. This page requires you to reset the default password for the admin login.

FIGURE 48 Setup Wizard: Set Administration Credentials Page

Pulse Secure

Set Administration Credentials

Change the default admin password to ensure that this system is secure. These are the credentials you should use to sign in to the administration web interface in future.

Username **admin**

Password

Confirm

NOTE

Changes made to the admin credentials will be applied immediately and will require user to authenticate using the new credentials.

[Previous](#) [Next](#)

6. Enter (and confirm) a password.

Note: The percent ("%") and UNICODE characters are not supported for this password.

Note: Administration credentials can be updated at any time after the Services Director VA is operational. See **"Updating Administration Credentials" on page 119**.

7. Click **Next**.

The Services Director VA login page appears.

8. Log into the Services Director VA using the new password.

On all platforms but AWS, the **Network Configuration** page appears.

Note: If your Services Director VA is on AWS, continue from step 11.

FIGURE 49 Setup Wizard: Network Configuration Page

Pulse Secure®

Network Configuration

Configure the network interface for this system.

Network Interface

☒ Apply Static IP

IP Address

Subnet Mask

Gateway

☐ Continue to use DHCP Allocated IP (not recommended)

NOTE

Changes made to the interface settings will be applied immediately and will require navigating back to this page using the new IP address.

Static Routes

+ Add

Destination	Subnet Mask	Gateway
No Data		

[Previous](#) [Apply](#)

9. Select one of the following options:

- **Static IP.** Then, complete an **IP Address** for the node (not the SEA), a **Subnet Mask** and a **Gateway**.

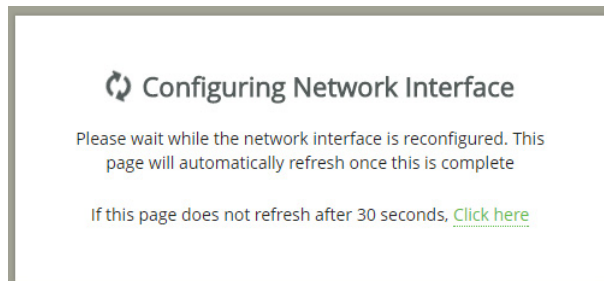
Note: The system will confirm that the gateway can be pinged.

- **DHCP Allocated IP.** Pulse Secure does not recommend the use of this option. A DHCP server must be available so that the system can request the IP address from it.

10. Click **Apply**.

A progress screen appears while the network interface is configured.

FIGURE 50 Setup Wizard: Configuring Network Interface Page

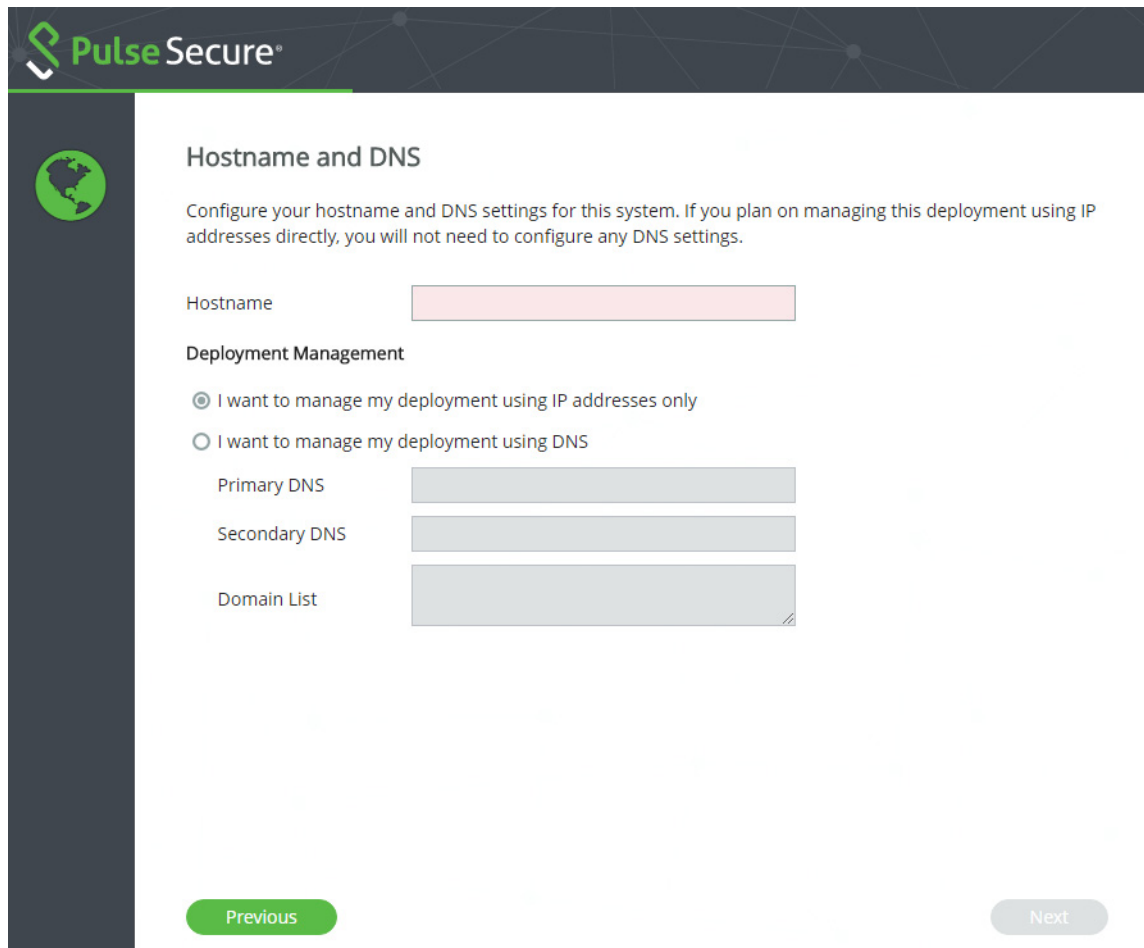


The outcome of this process depends on whether you selected **Static IP** or **DHCP Allocated IP**.

- **Static IP.** The browser will automatically access the wizard using the specified IP address. Log in, and continue the Setup Wizard.
- **DHCP Allocated IP.** Manually direct your browser to the allocated IP to continue this wizard. Log in, and continue the Setup Wizard.

The **Hostname and DNS** page appears. This page enables you to choose whether to manage your Services Director using either IP addresses or DNS.

FIGURE 51 Setup Wizard: Hostname and DNS Page



The screenshot shows the 'Hostname and DNS' configuration page in the Pulse Secure Setup Wizard. The page has a dark header with the Pulse Secure logo and a green sidebar on the left containing a globe icon. The main content area is white and contains the following elements:

- Hostname and DNS** title.
- Instructional text: "Configure your hostname and DNS settings for this system. If you plan on managing this deployment using IP addresses directly, you will not need to configure any DNS settings."
- A text input field for the **Hostname**.
- A section titled **Deployment Management** with two radio button options:
 - ☒ I want to manage my deployment using IP addresses only
 - ☐ I want to manage my deployment using DNS
- Below the DNS option, there are three text input fields for **Primary DNS**, **Secondary DNS**, and **Domain List**.
- At the bottom, there are two buttons: a green **Previous** button and a grey **Next** button.

11. On the **Hostname and DNS** page, enter the management address for the Services Director as the **Hostname**.

- If this management address can be resolved using DNS, enter its hostname.
- If this management address cannot be resolved using DNS, enter its IP address.

Where no DNS is configured, the use of hostnames should be avoided in the product.

12. Select one of the following options:

- **I want to manage my deployment using IP addresses only.** Select this where no DNS is configured.

Ensure that you specify the Services Director's IP address as its **Hostname** (see above).

- **I want to manage my deployment using DNS.** This requires you to have one or more configured DNS name servers in place.

Ensure that you specify a resolvable hostname as the Services Director's **Hostname** (see above). Then, specify:

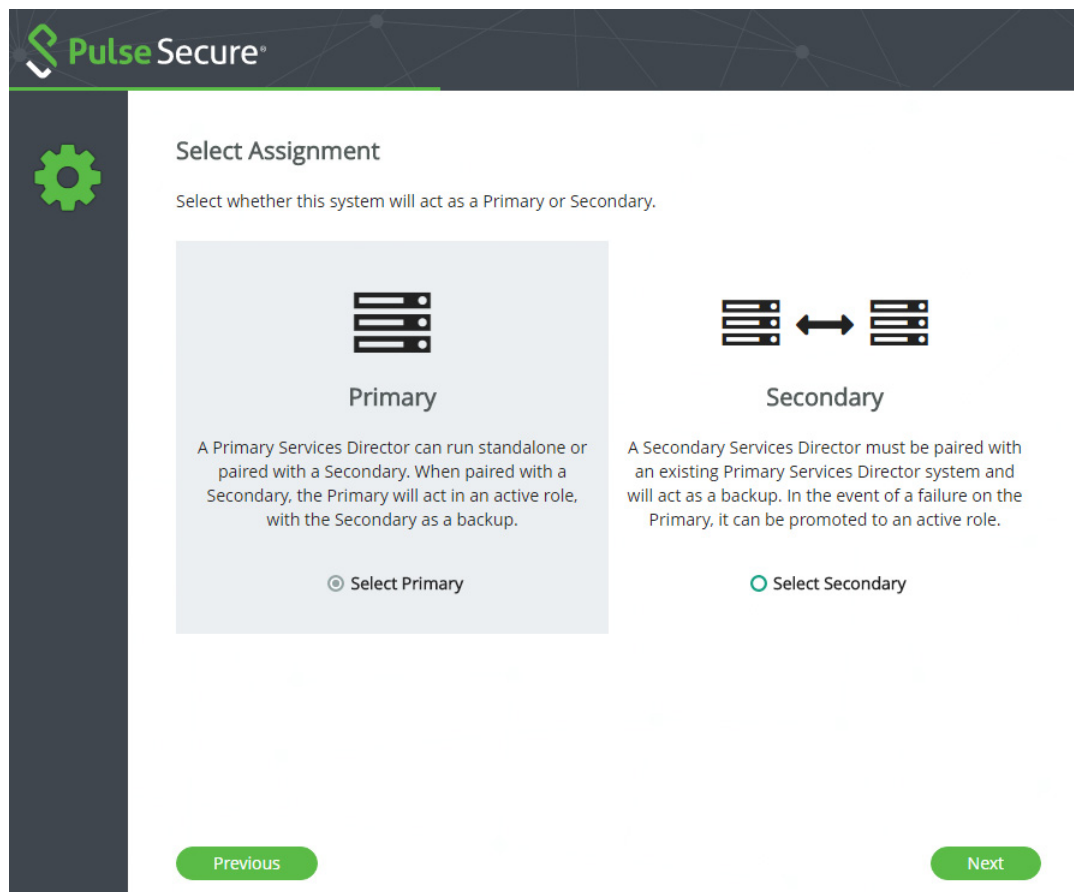
- **Primary DNS**
- **Secondary DNS** (Optional)
- **Domain List** (Optional) An ordered list of domain names. If you specify domains, the system automatically finds the appropriate domain for each of the hosts that you specify in the system.

13. Click **Next**.

The **Select Assignment** page appears.

This page enables you to select the role of the Services Director.

FIGURE 52 Setup Wizard: Select Assignment Page



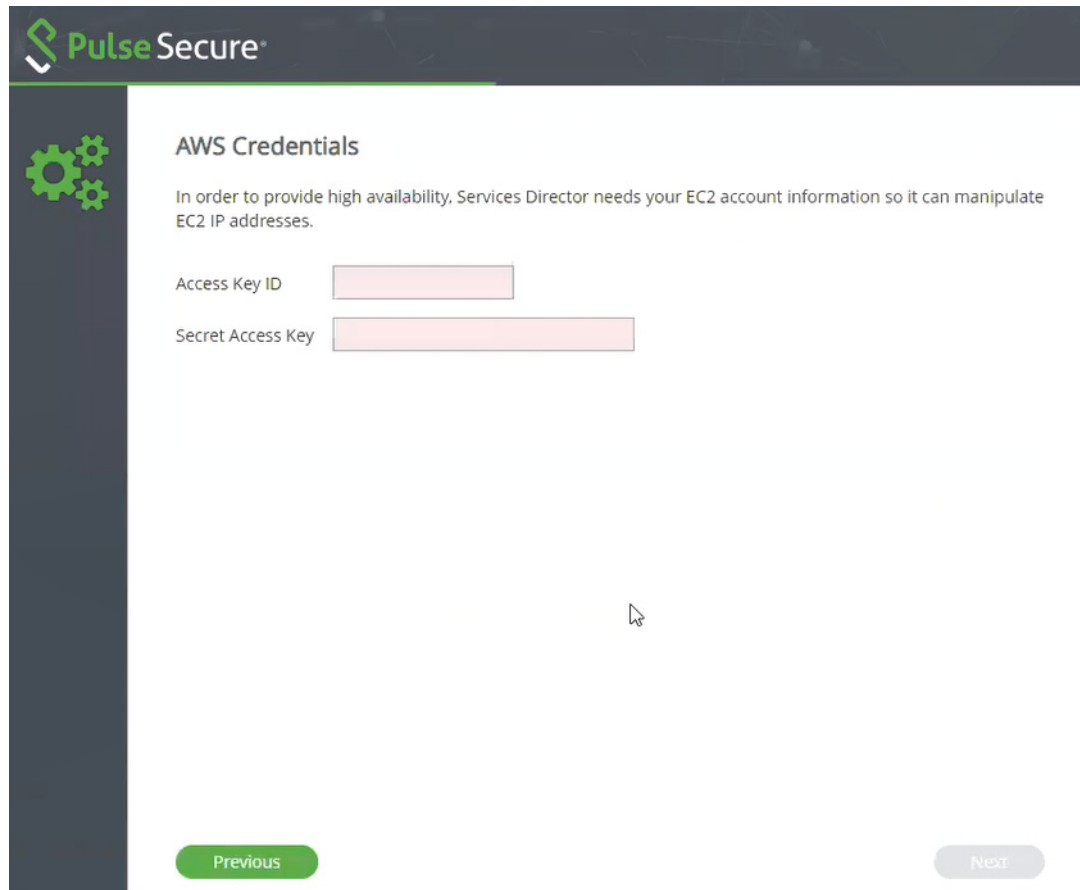
14. Click **Select Primary** to indicate that the Services Director will act as a Primary Services Director, either as a standalone node or in an HA Pair.
15. Click **Next**.

For vSphere and KVM hypervisor platforms, continue from this process from **“Defining a Service Endpoint Address” on page 89**.

For the AWS platform only:

- Where you specified an IAM Role during the launch of the Services Director instance, continue from this process from **“Defining a Service Endpoint Address” on page 89**.
- Where you did not specify an IAM Role, the **AWS Credentials** page appears:

FIGURE 53 Setup Wizard: AWS Credentials



The screenshot shows the 'AWS Credentials' page of the Pulse Secure Setup Wizard. The page has a dark header with the 'Pulse Secure' logo. On the left, there is a dark sidebar with a green gear icon. The main content area is white and contains the following text: 'AWS Credentials', 'In order to provide high availability, Services Director needs your EC2 account information so it can manipulate EC2 IP addresses.', 'Access Key ID' with a text input field, and 'Secret Access Key' with a text input field. At the bottom, there are two buttons: a green 'Previous' button and a grey 'Next' button.

16. Enter your AWS credentials.

17. Click **Next**, and continue this process from **“Defining a Service Endpoint Address” on page 89**.

Defining a Service Endpoint Address

After you choose to define a Primary Services Director, the **Service Endpoint Address** page appears.

FIGURE 54 Setup Wizard: Define a Service Endpoint Address Page

Pulse Secure

Service Endpoint Address

Choose a Service Endpoint IP address that will be used by this system. The Service Endpoint IP is used to ensure high-availability as in the event of a failover, the Secondary Services director will be available via the same IP address that the Primary was accessible from.

Service Endpoint IP Address

☒ The Service Endpoint Address is globally addressable
☐ The Service Endpoint Address is behind a NAT device

NOTE

After Setup is complete, you should use the **Service Endpoint Address** to locate this system, not the IP used by the network interface in the Network Configuration step. This is also the IP address you should provide to Pulse Secure in order to generate your FLA license (or if you supply a hostname, a hostname which maps to this IP address).

Previous Next

1. If the Service Endpoint Address (SEA) for the Services Director HA pair will be routed to directly by the vTMs in its estate:
 - Select **The Service Endpoint Address is globally addressable**.
 - Enter the required **Service Endpoint IP Address** for the Services Director HA pair.

A Service Endpoint Address is required for a standalone Primary Services Director. It must be different from the IP address of the Primary Services Director.

2. If the SEA for the Services Director HA pair is behind a NAT device (from the point of view of the vTMs that will be in its estate):
 - Select **The Service Endpoint Address is behind a NAT device**. The available properties update to include an **External IP Address**.
 - Enter the internal NAT SEA for your Services Director HA pair as the **Service Endpoint IP Address**.

- Enter the external NAT address for your Services Director HA pair as the **External IP Address**.

Note: A Service Endpoint Address is required for a standalone Primary Services Director. It must be different from the IP address of the Primary Services Director.

3. Click **Next**.

The **Restore from Backup** page appears.

This page enables you to restore a backup of your Services Director after a failure. Refer to the *Pulse Services Director Advanced User Guide* for details.

FIGURE 55 Setup Wizard: Restore from Backup Page

Pulse Secure

Restore from backup

If you have a backup file from a previous installation, you can restore it now. Otherwise you can proceed with a new installation.

☒ This is a new system

☐ Restore from a previous backup

[Choose File](#)

Master Password: ☒ Save the password?

NOTE

Backup files do not include vTM images that may have been in use. If you were using managed vTM instances in your previous installation, you will need to re-upload the vTM image files separately after completing Setup.

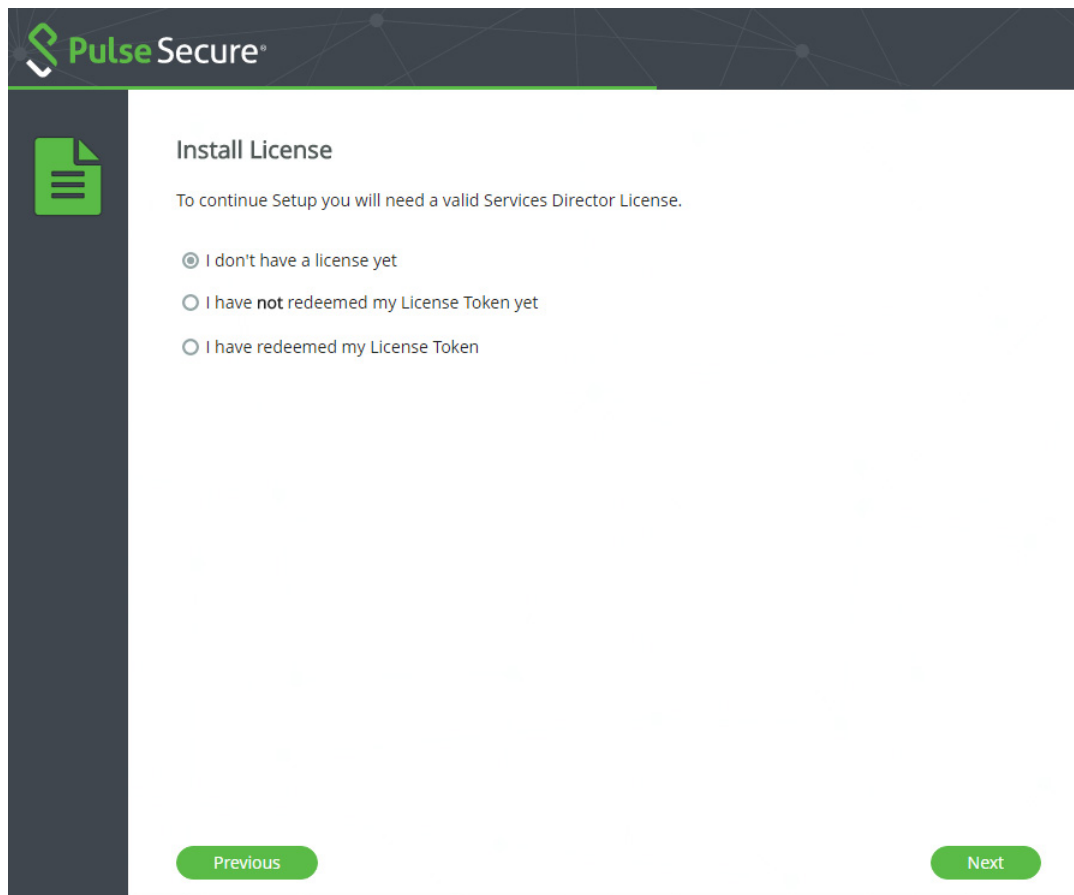
For security, it is recommended that the master password is input manually every time the Services Director starts. However, the password could be stored in a file (less secure) for non-interactive start up.

[Previous](#) [Next](#)

4. Click **This is a new appliance** and then click **Next**.

The **Install License** page appears.

FIGURE 56 Setup Wizard: Install License Page



5. Select one of the following options:
 - **I have redeemed my License Token.** You can now add your licenses. Click **Next**, and continue from **"Adding Certificates and Licenses" on page 95**.
 - **I have not redeemed my License Token yet.** The Setup Wizard will guide you through this process. Click **Next**, and continue from **"Redeeming a License Token" on page 92**.
 - **I don't have a license yet.** If you have not obtained a License Token, you *cannot* proceed with the Setup Wizard at this time. See **"Obtaining Services Director Licenses" on page 14**.

Close the Setup Wizard.

Redeeming a License Token

After you indicate that you have an unredeemed license token, the **SSL Certificate Generate** page appears. An SSL certificate is required to redeem your token. You can provide your own certificate, or the system can generate one for you.

FIGURE 57 Setup Wizard: SSL Certificate Generate Page

Pulse Secure

SSL Certificate Generate

To redeem your Token, you must provide Pulse Secure with a self-signed SSL Certificate. This SSL certificate is only used to secure the licensing system. A self-signed certificate can be generated for you, or you can provide your own self-signed certificate.

☒ Generate a self-signed certificate for me

☐ I will provide my own self-signed certificate (in PEM format)

NOTE

You should NOT provide a CA-signed SSL certificate, if you wish to generate a self-signed certificate in PEM format using OpenSSL, you can do so with the following command:

```
openssl req -x509 -nodes -newkey rsa:1024 -keyout key.pem -out cert.pem -days 3650
```

[Previous](#) [Next](#)

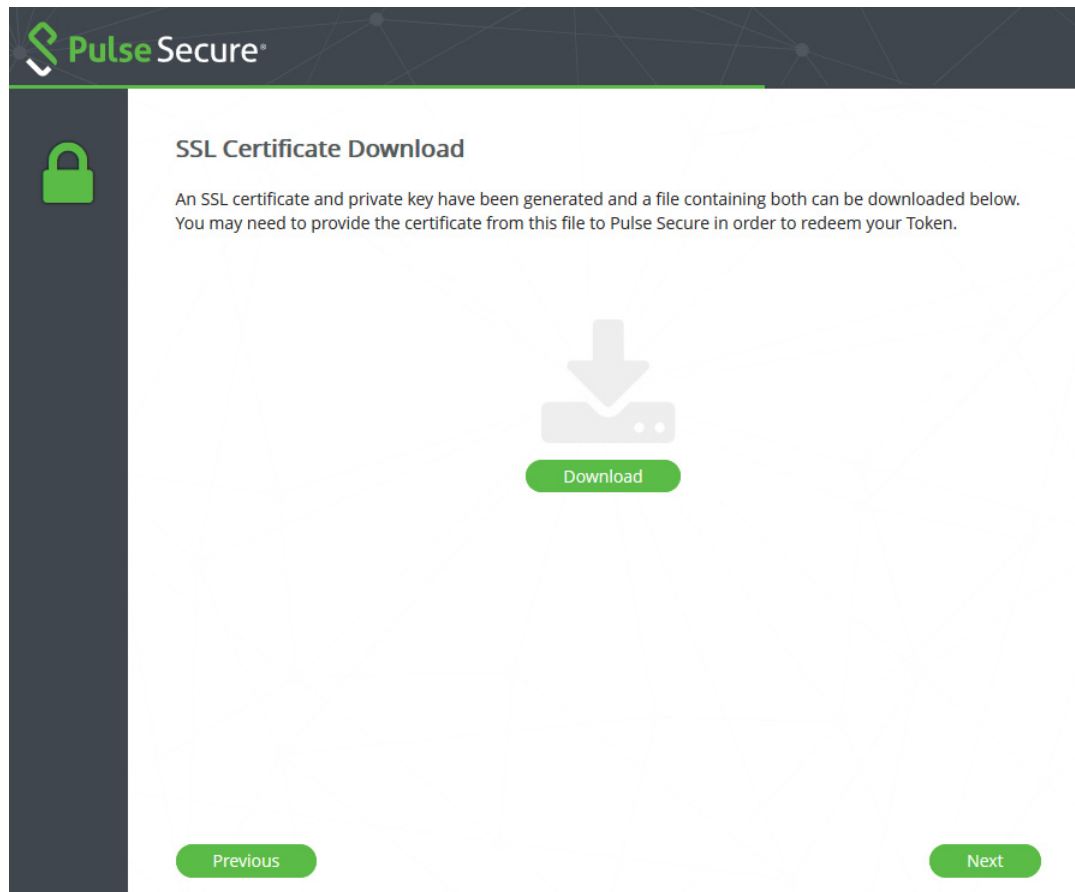
Select one of the following options:

- **Generate a signed certificate for me.** This selection will instruct the system to create a signed certificate that can be used to redeem your License Token with Pulse Secure. Click **Next**, and continue from **"Generating a Self-Signed SSL Certificate" on page 93**.
- **I will provide my own self-signed certificate.** This selection requires you to have a self-signed SSL certificate. *You cannot use a CA-signed certificate.* Click **Next**, and continue from **"Adding Certificates and Licenses" on page 95**.

Generating a Self-Signed SSL Certificate

After you choose to have Services Director generate a self-signed SSL certificate, the **SSL Certificate Download** page appears. An SSL certificate is required to redeem your token.

FIGURE 58 Setup Wizard: SSL Certificate Download Page

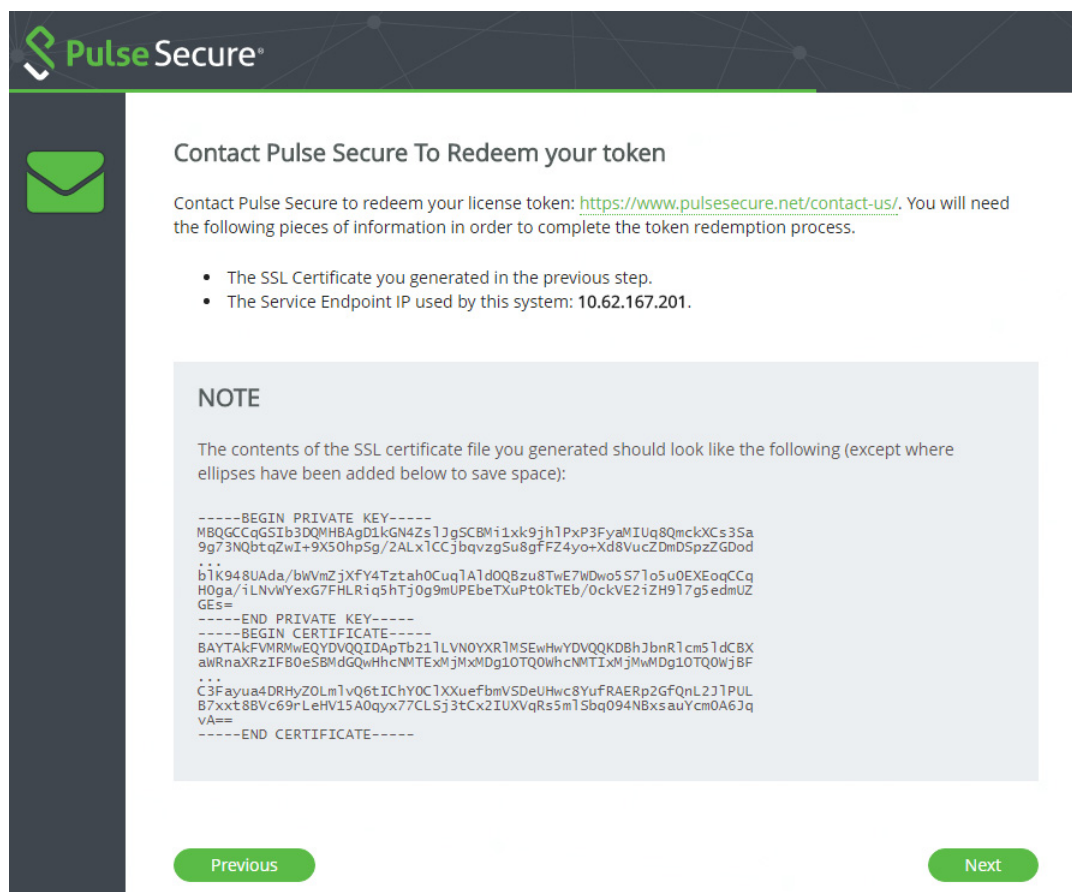


1. Click **Download** and choose a location for the file. The self-signed SSL certificate file downloads.
2. Click **Next**.

The **Contact Pulse Secure to Redeem Your Token** page appears. This page provides advice about how to redeem your token.

Note: You cannot proceed with the Setup Wizard until you have redeemed your token.

FIGURE 59 Setup Wizard: Contact Pulse Secure to Redeem Your Token Page



3. To redeem your License Token, visit the Pulse Secure License Redemption Portal.

- Your License Token.
- Your self-generated SSL certificate.
- The Service Endpoint Address.

Once you have your licenses, continue from **“Adding Certificates and Licenses” on page 95**.

Adding Certificates and Licenses

After you have redeemed your License Token, the **SSL Certificate Upload** page appears. This page enables you to input your certificate. The text of the certificate can be pasted in manually. Alternatively, you can identify individual Private/public key files, or a single combined file.

Note: If you previously chose to generate a self-signed certificate using the Setup Wizard, you will bypass this screen. This is because the Services Director already has the SSL certificate.

FIGURE 60 Setup Wizard: SSL Certificate Upload Page

Pulse Secure

SSL Certificate Upload

Upload the SSL certificate and private key you provided to Pulse Secure when you redeemed your license Token.

☒ Single file with public and private keys

[Choose File](#)

☐ Separate public and private key files

Private key [Choose File](#)

Public key [Choose File](#)

☐ Text content of the public and private keys

[Previous](#) [Next](#)

1. Select one of the following options:

- **Single file with public and private keys.** Then, click **Choose File** to locate the certificate file.
- **Separate public and private key files.** Then, click **Choose File** to locate each file.
- **Text content of the public and private keys.** Then, paste the required text in.

The selected text/file(s) are then verified. If successful, the **Next** button becomes available.

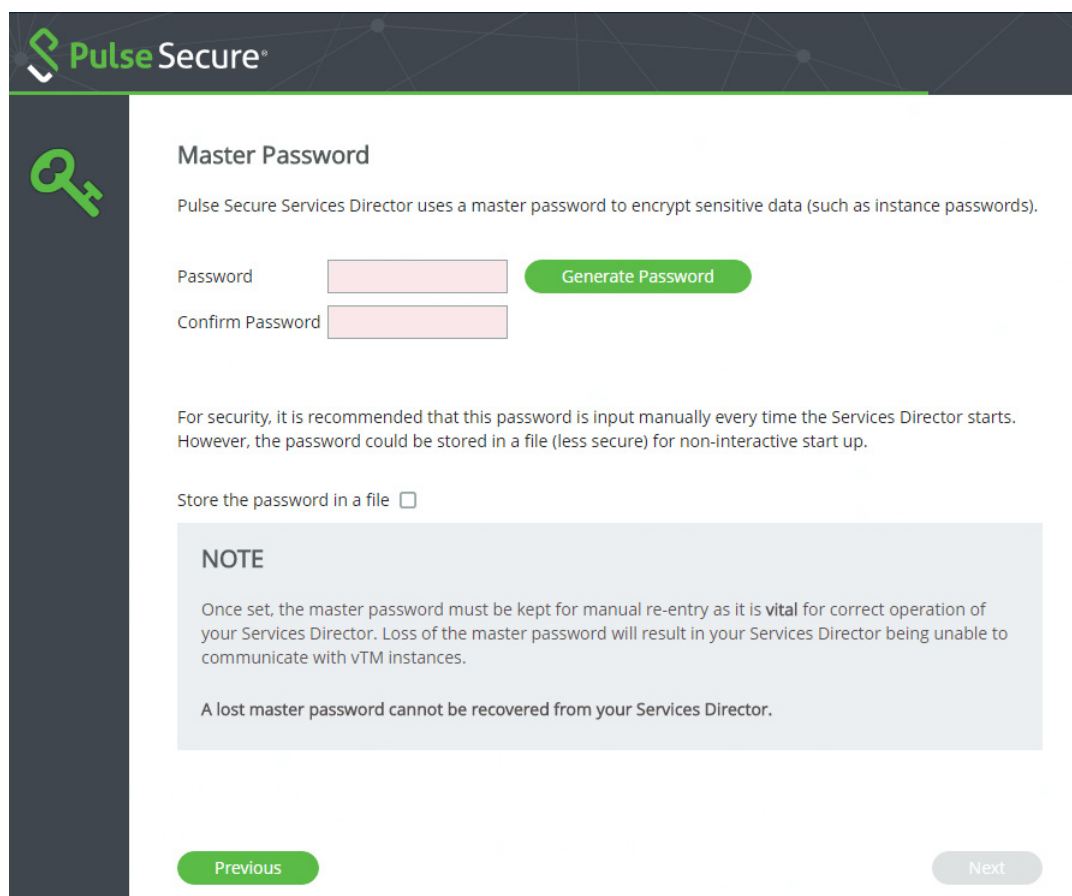
The SSL certificate can be changed after the Services Director VA is operational. See **“Updating the SSL Certificate” on page 119.**

2. Click **Next**.

The Services Director **Master Password** page appears. This page enables you to define a master password. A master password is required to:

- To decrypt stored password information whenever the Virtual Machine for this Services Director VA node restarts.
- To create a new Services Director VA from a previously-saved backup, see [“Recovering from a Services Director Failure” on page 401](#)).

FIGURE 61 Setup Wizard: Master Password Page



The screenshot shows the 'Master Password' page in the Pulse Secure setup wizard. The page has a dark header with the Pulse Secure logo and a dark sidebar on the left with a green key icon. The main content area is white. It features a title 'Master Password', a description of its purpose, two password input fields with a 'Generate Password' button, a checkbox for storing the password in a file, a 'NOTE' box with important warnings, and 'Previous' and 'Next' navigation buttons at the bottom.

Pulse Secure

Master Password

Pulse Secure Services Director uses a master password to encrypt sensitive data (such as instance passwords).

Password Generate Password

Confirm Password

For security, it is recommended that this password is input manually every time the Services Director starts. However, the password could be stored in a file (less secure) for non-interactive start up.

Store the password in a file ☐

NOTE

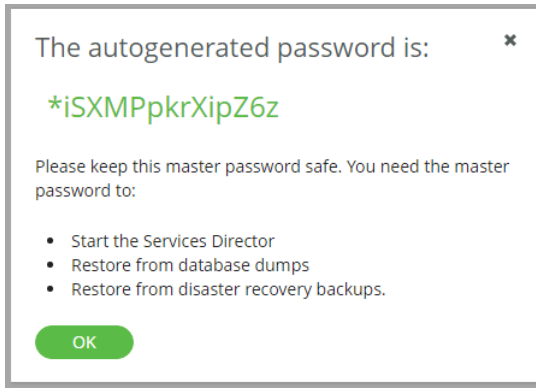
Once set, the master password must be kept for manual re-entry as it is **vital** for correct operation of your Services Director. Loss of the master password will result in your Services Director being unable to communicate with vTM instances.

A lost master password cannot be recovered from your Services Director.

Previous Next

3. To set the master password, perform one of the following operations.
 - Enter a password and confirm the password.
 - Click **Generate Password**. The **Password** and **Confirm Password** fields are populated automatically and a dialog box is displayed.

FIGURE 62 Setup Wizard: Master Password Dialog



Record the password, click **OK** to close the information dialog box, and then confirm that you have stored the password in the next dialog box.

Note: It is essential that the master password (whether chosen yourself or generated automatically) is recorded and can be retrieved. Pulse Secure recommends that this password is recorded in a secure location that is separate from the Services Director VA.

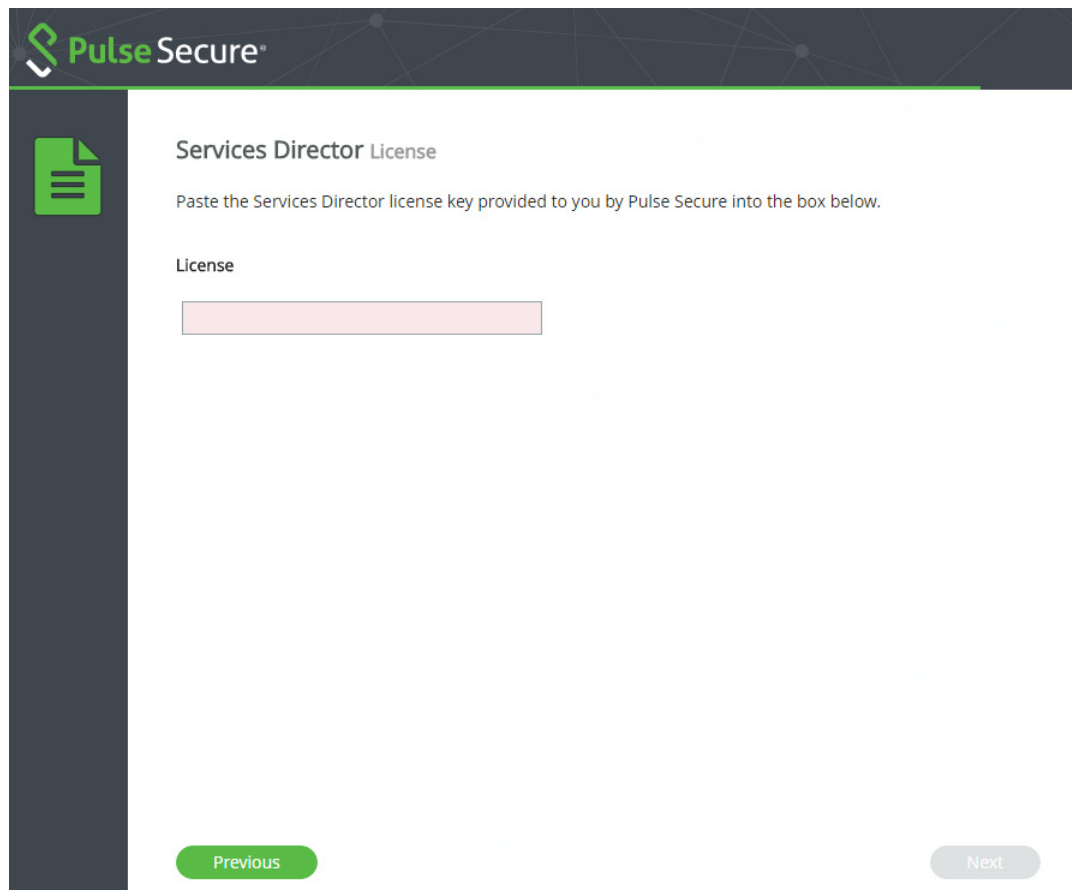
4. Choose whether to store the password internally for automatic use:
 - Select the **Store the password in a file** check box to store the master password within the Services Director VA. The password will be automatically available whenever the Virtual Machine for a Services Director VA restarts. However, you must enter the master password manually when you create a Services Director VA from a backup file.
 - Clear the **Store the password in a file** check box to not record the master password. You must to enter the master password manually whenever the Virtual Machine for a Services Director VA restarts, and when you create a Services Director VA from a backup file.

See **“Entering the Master Password After a Virtual Machine Restart” on page 418** for details of restarting a VM.

5. Click **Next**.

The Services Director **License** page appears.

FIGURE 63 Setup Wizard: Services Director License Page



The screenshot shows the 'Services Director License' page of the Pulse Secure Setup Wizard. The page has a dark header with the Pulse Secure logo. A green sidebar on the left contains a document icon. The main content area is white and contains the title 'Services Director License', a instruction to paste the license key, a text input field labeled 'License', and two buttons at the bottom: 'Previous' (green) and 'Next' (gray).

Pulse Secure®

Services Director License

Paste the Services Director license key provided to you by Pulse Secure into the box below.

License

Previous Next

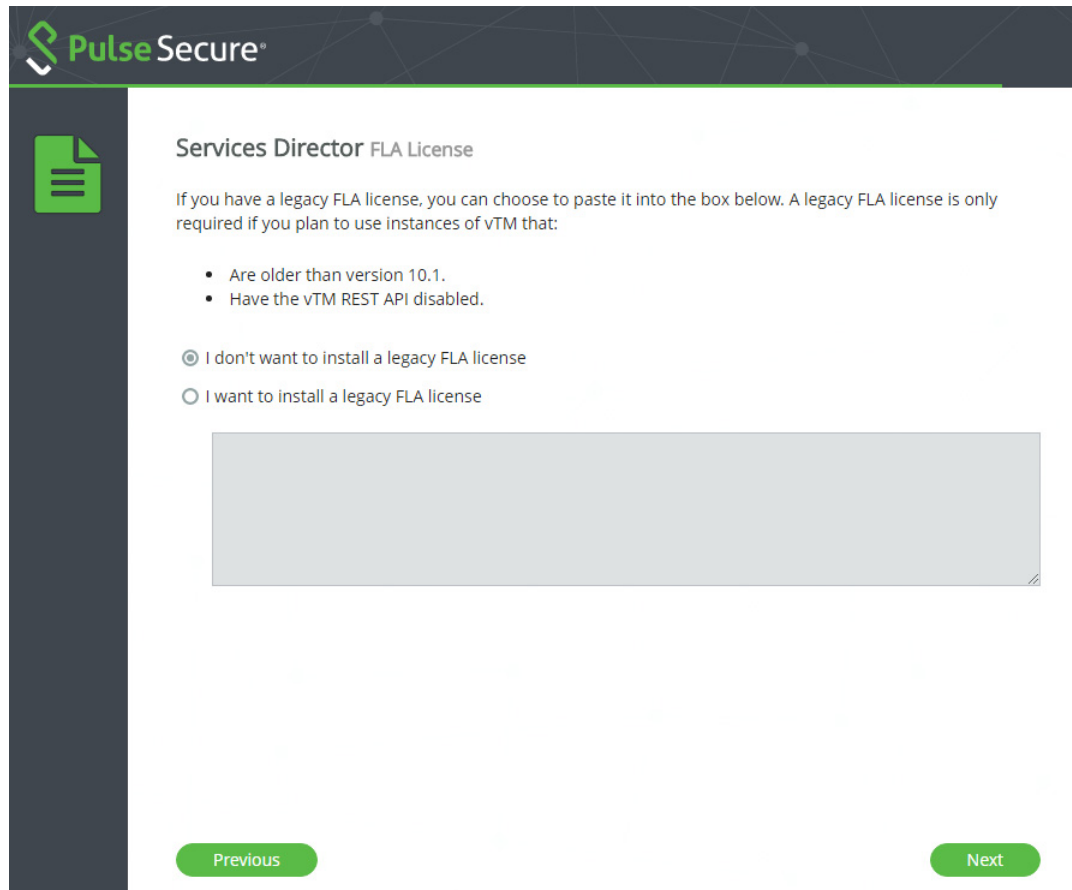
6. Enter the **License** text. This is validated automatically. Once validation completes, either:
 - A success message is displayed, and the **Next** button becomes available. OR
 - A failure message is displayed. You must repeat this step.

7. Click **Next**.

The Services Director **FLA License** page appears.

This page enables you to add a Legacy FLA license if you are using a vTM at version 10.0 (or earlier), or wish to disable the REST API for any of your vTM instances.

FIGURE 64 Setup Wizard: Services Director FLA License Page

The screenshot shows the 'Services Director FLA License' page within the Pulse Secure Setup Wizard. The page has a dark header with the Pulse Secure logo. A green sidebar on the left contains a document icon. The main content area is white and contains the following text: 'Services Director FLA License', 'If you have a legacy FLA license, you can choose to paste it into the box below. A legacy FLA license is only required if you plan to use instances of vTM that:', a bulleted list of conditions (older than version 10.1, REST API disabled), and two radio button options: 'I don't want to install a legacy FLA license' (selected) and 'I want to install a legacy FLA license'. Below the options is a large gray text box for pasting a license. At the bottom are 'Previous' and 'Next' buttons.

8. Select one of the following options:

- **I don't want to install a legacy FLA license.** You will do this for one of the following reasons:
 - You want to use the installed Universal FLA License. To support this selection, all of your vTM instances must be running version 10.1 (or later) with the REST API enabled.
 - You do not want to install a Legacy FLA License at this time. This can be entered using the Services Director VA graphical interface after it is deployed.

A default Feature Pack will not be created, but this can be created at a later date. See **"Adding a Feature Pack to the Services Director" on page 129**.

Continue from the next step.

- **I want to install a legacy FLA license.** You will do this if any of your vTMs are running at version 10.0 (or earlier) or have their REST API disabled. Paste the text of your Legacy FLA License into the box. This is validated automatically.

9. Click **Next**.

The Services Director **Additional Licenses** page appears.

- If you have and Resource Licenses, either for bandwidth or analytics, use this page to enter them.
- If you do *not* have Resource Licenses at this point, you can still continue with the Setup Wizard. You can enter these licenses using the Services Director VA after it is deployed.
- If you have a Cloud Services Provider (CSP) License for your Services Director, you do not require Resource Licenses, and can ignore this page.

FIGURE 65 Setup Wizard: Services Director Additional Licenses Page

Services Director Additional licenses

Add in any additional licenses that you have been provided by Pulse Secure. These licenses may consist of one or more Bandwidth, Resource, or Add-on licenses. If you do not add these licenses now, you can do so after completing Setup by going to the licenses page in the admin web interface.

Add license

Add

Additional licenses

License	Type
No Data	

Previous Next

10. Enter a license number and click **Add**.

This license is validated automatically. Once validation completes, the license is listed in the **Additional licenses** table, along with its type.

11. Repeat the previous step to add all available licenses.

12. Click **Next**.

The **Email alerts** page appears.

This page enables you to optionally enter email notification details for your Services Director. This ensures that you receive email notifications for events and failures.

Note: You do not have to enter this information now. It can be entered using the Services Director VA after it is deployed. See [“Updating Email Settings” on page 119](#).

FIGURE 66 Setup Wizard: Email Alerts Page

Pulse Secure

Email Alerts

Configure your SMTP settings to enable this system to send email alerts. It is highly recommended that you set up email alerts as this is the only notification mechanism available to inform you of problems with this system.

☒ I do not want to configure email alerts

☐ I want to configure email alerts (**Recommended**)

Destination email address

SMTP server

SMTP port

Telemetry

Services Director can collect and export anonymized usage information to Pulse Secure, to help improve our products. See this [Knowledge Base article](#) for details of the collected data.

☒ I want to enable telemetry (**Recommended**)

☐ I do not want to enable telemetry

13. Under **Email Alerts**, select one of the following options:

- **I do not want to configure email alerts.** This option enables you to bypass this step. This information can be entered using the Services Director VA graphical interface after it is deployed. See [“Updating Email Settings” on page 119](#).

- **I want to configure email alerts.** This is the recommended option. Then, provide:
 - A **Destination email address**.
 - An **SMTP server**. This is either the hostname or IP address of the SMTP server in your network.
 - An **SMTP port** number. Typically, you will use the default port number, 25.

14. Click **Send test email** to confirm these settings.

Note: You must have external access for SMTP traffic for this feature to function.

15. Under **Telemetry**, select whether you want Services Director to collect and export anonymized usage information to Pulse Secure.

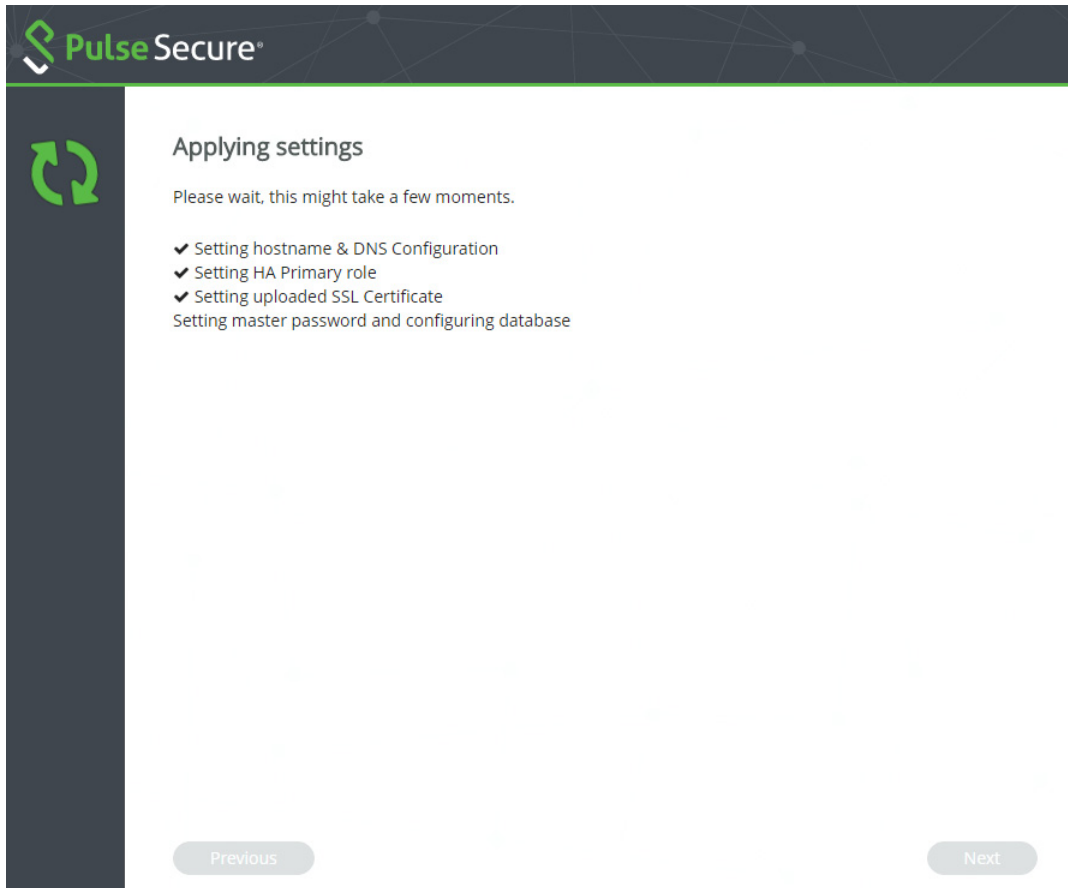
Note: This setting can be changed from the **General Settings** page at any time, see [“Updating Telemetry Settings” on page 117](#).

16. Click **Next**, and continue from [“Completing the Services Director Installation” on page 103](#).

Completing the Services Director Installation

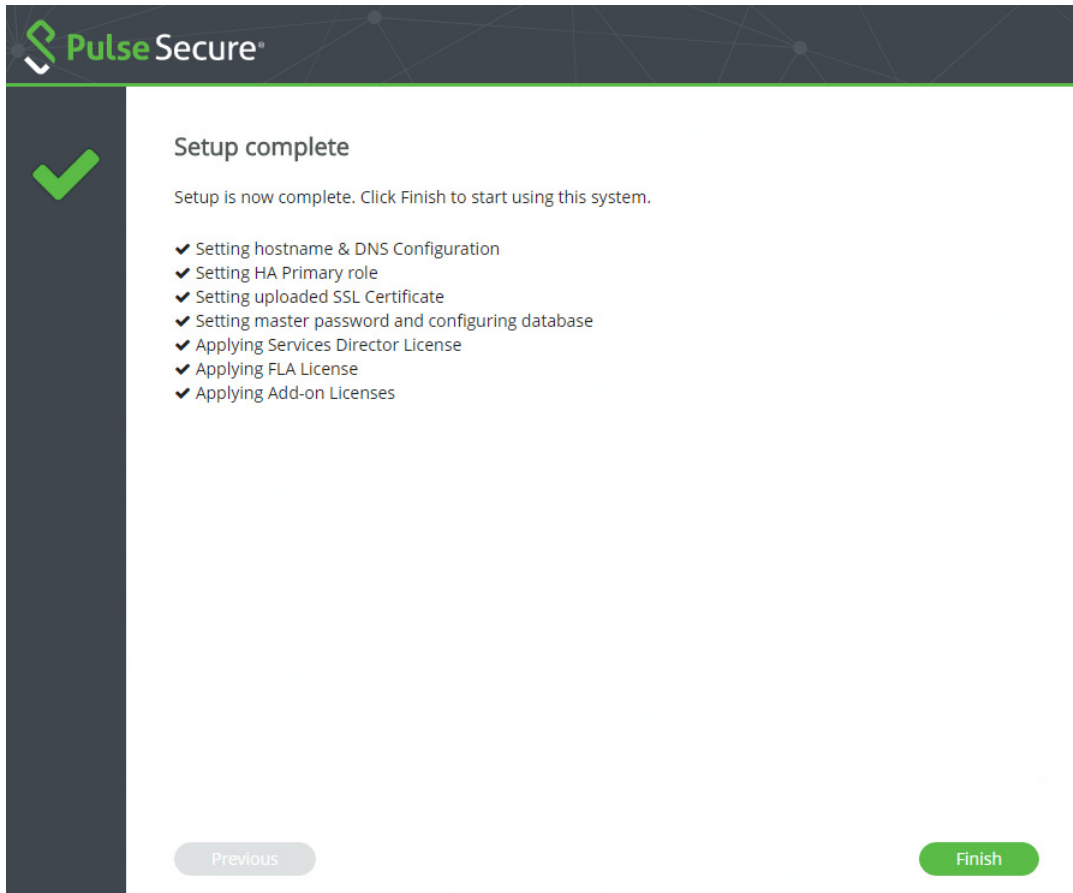
After all information is gathered, the **Applying Settings** page appears. This page configures the system based on collected information. For example:

FIGURE 67 Setup Wizard: Applying Settings Page



Once this is complete, the **Setup Complete** page appears.

FIGURE 68 Setup Wizard: Setup Complete Page



1. Click **Finish** to close the Setup Wizard.

Once the Setup Wizard completes, your Services Director node is ready for use.

2. (Optional) you can now create a Secondary Services Director, and join it to the Primary Services Director. See **"Installing and Configuring a Secondary Services Director" on page 105.**

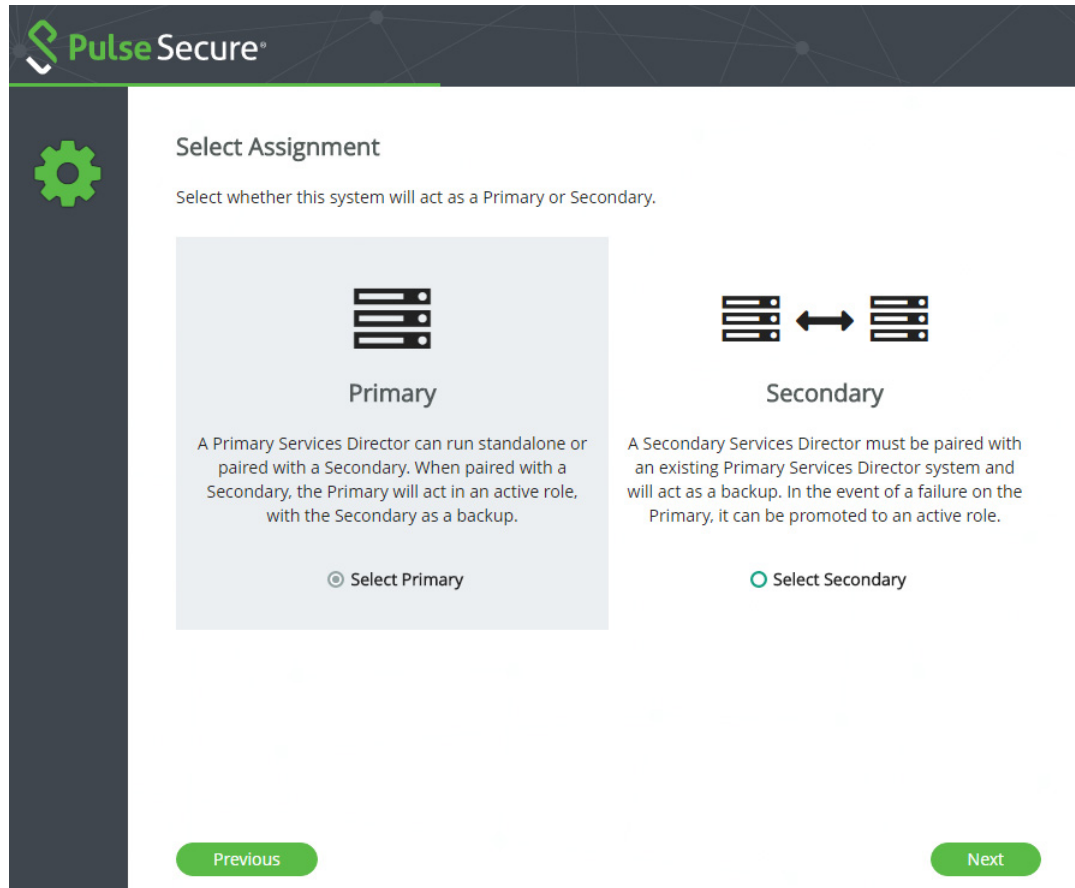
Note: Once the Setup Wizard completes, it cannot be rerun. Many of the options chosen in the Setup Wizard can be reconfigured from inside the Services Director VA, but others can only be reconfigured from the Command-Line Interface (CLI). See *Pulse Services Director Advanced User Guide* and the *Secure Services Director Command Reference* for full details.

Installing and Configuring a Secondary Services Director

The process for creating a Secondary Services Director is similar to the installation for a Primary Services Director.

1. Repeat the installation process for a Primary Services Director (see [“Starting the Setup Wizard” on page 80](#)) until you reach the following screen:

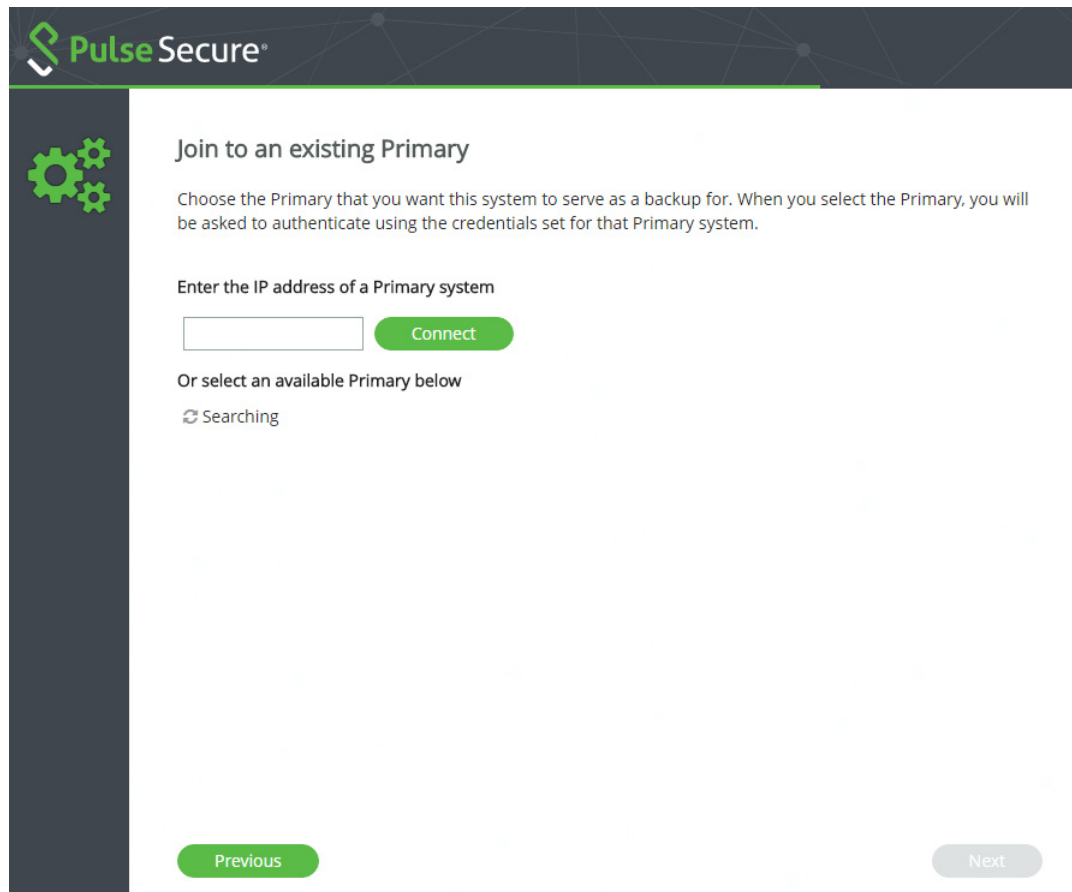
FIGURE 69 Setup Wizard: Select Assignment



2. Click **Select Secondary**.

The **Join to an Existing Primary** page appears.

FIGURE 70 Setup Wizard: Join to an Existing Primary Page



Pulse Secure®

Join to an existing Primary

Choose the Primary that you want this system to serve as a backup for. When you select the Primary, you will be asked to authenticate using the credentials set for that Primary system.

Enter the IP address of a Primary system

Connect

Or select an available Primary below

Searching

Previous Next

3. To connect to an existing Primary Services Director, either:

- Select the Primary Services Director from the list.

Note: This option is not supported by the AWS platform.

- Enter the IP address of the Primary Services Director.

Note: On the AWS platform, this must be the Primary Private IP Address of the instance.

4. Click **Connect**.

The page updates to include an **Enter Credentials** panel.

FIGURE 71 Setup Wizard: Join to an Existing Primary: Enter Credentials

The screenshot shows the Pulse Secure Setup Wizard interface. The header features the Pulse Secure logo. A sidebar on the left contains a gear icon. The main content area is titled "Join to an existing Primary" and includes instructions: "Choose the Primary that you want this system to serve as a backup for. When you select the Primary, you will be asked to authenticate using the credentials set for that Primary system."

There are two main sections for selection:

- Enter the IP address of a Primary system:** A text input field followed by a green "Connect" button.
- Or select an available Primary below:** A list of available primaries:
 - gold-01 >
 - rmccann-04 >
 - rmccann-02 >
 - amnesiac >
 - dmankellow-3b >

To the right, the **Enter credentials** panel is shown for "gold-01". It includes fields for "Username" and "Password", and an "Authenticate" button.

At the bottom, there are "Previous" and "Next" buttons.

5. Under **Enter credentials**, enter an administration login details for the Primary Services Director.

6. Click **Authenticate**.

The credentials are confirmed.

7. Click **Next**.

The Services Director **Master Password** page appears.

This page requires you to enter the master password that you chose for the Primary Services Director VA. This is required to:

- To decrypt stored password information whenever the Virtual Machine for this Services Director VA node restarts.

- To create a new Services Director VA from a previously-saved backup, see [“Recovering from a Services Director Failure” on page 401](#)).

FIGURE 72 Setup Wizard: Services Director Master Password Page

Master Password

This cluster uses a master password to encrypt sensitive information. Please enter the master password as configured on the Primary HA node.

Password

For security, it is recommended that this password is input manually every time the Services Director starts. However, the password could be stored in a file (less secure) for non-interactive start up.

Store the password in a file ☐

[Previous](#) [Next](#)

8. Enter the master password. The password is validated immediately.
9. Choose whether to store the password internally for automatic use:
 - Select the **Store the password in a file** check box to store the master password within the Services Director VA. The password will be automatically available whenever the Virtual Machine for a Services Director VA restarts. However, you must enter the master password manually when you create a Services Director VA from a backup file.
 - Clear the **Store the password in a file** check box to not record the master password. You must to enter the master password manually whenever the Virtual Machine for a Services Director VA restarts, and when you create a Services Director VA from a backup file.

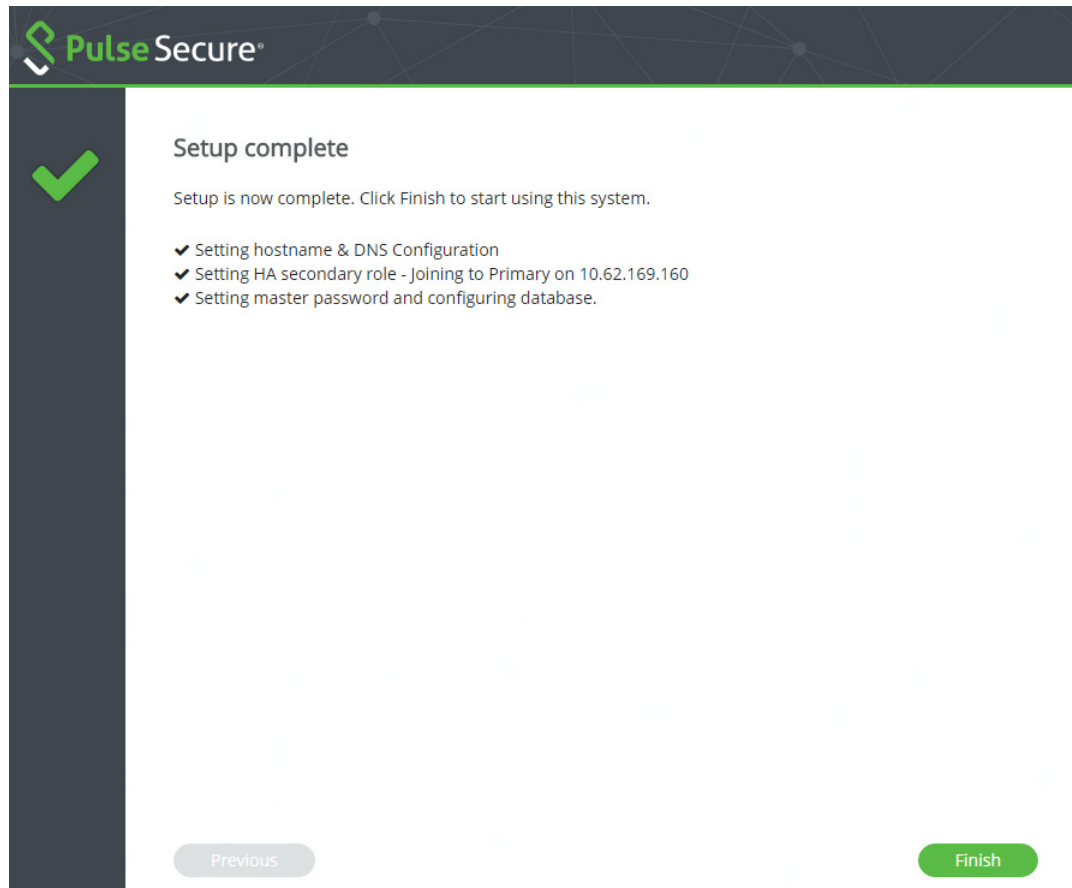
See [“Entering the Master Password After a Virtual Machine Restart” on page 418](#) for details of restarting a VM.

10. Click **Next**.

The Secondary Services Director now joins with the Primary Services Director to form a HA pair. The progress of this process appears on the **Applying Settings** page.

Once this process completes, the **Setup Complete** page appears.

FIGURE 73 Setup Wizard: Setup Complete Page (Secondary Services Director)



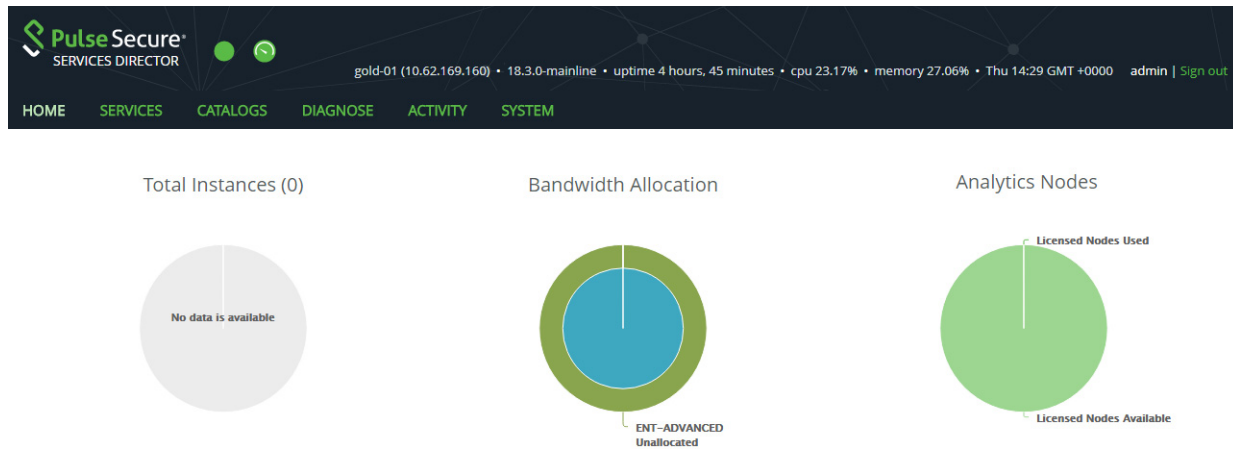
Accessing your Services Director VA

Once the Setup Wizard is complete, you can access the Services Director VA using a secure (https) URL in a browser:

- For an HA pair, you access the *Active* Services Director using the Service Endpoint IP address.
If the Services Director HA pair is in a private network behind a NAT device, access the *Active* Services Director using the external IP address of the Service Endpoint Address.
- You can access a standalone Services Director using its IP address or Service Endpoint IP address.
- You can access the Primary Services Director directly using its IP address.
- You can access the Secondary Services Director directly using its IP address.

Log in to the Services Director VA. The **Home** page appears:

FIGURE 74 The Home Page



The header displays two coloured indicators:

- The first is an indication of system health. This includes: high availability, the Services Director license, and the availability of the service.
- A healthy system displays a green circle, and an unhealthy system displays an orange warning triangle.
- The second is an indicator for metering discrepancies for the vTMs within the estate of the Services Director VA.

A healthy metering system results in a green meter. An unhealthy metering system displays as an orange warning meter. See **“Processing Virtual Traffic Manager Metering Discrepancy Warnings” on page 224.**

At this point, no vTMs are registered on the Services Director VA.

The **Home** page always displays:

- The **Total Instances** of vTM vTMs registered on the Services Director.

Note: Immediately after the Services Director is installed, there are zero registered vTMs.

- The **Bandwidth Allocation** for all Bandwidth Licenses that were installed during the Setup Wizard.

Note: Immediately after the Services Director is installed, there are zero allocations.

- The **Analytics Nodes** for all Analytics Resource Pack Licenses that were installed during the Setup Wizard.

Note: Immediately after the Services Director is installed, there are zero licensed nodes.

Optionally, you may wish to fine-tune settings for the Services Director VA. See [“Updating Services Director VA Settings” on page 113](#).

Otherwise, you can now proceed with the registration of vTMs and additional system configuration. See [“Adding Virtual Traffic Managers to the Services Director” on page 125](#).

Updating Services Director VA Settings

- **Overview: Services Director VA Settings** 113
- **Updating General Settings**..... 113
- **Updating Date and Time Settings**..... 118
- **Updating Administration Credentials** 119
- **Updating Email Settings**..... 119
- **Updating the SSL Certificate** 119
- **Updating the REST API Port**..... 120
- **Updating Security Settings** 120
- **Changing the Master Password for the Services Director VA** 121

Overview: Services Director VA Settings

Once your Services Director VA is installed on your chosen platform, you can configure the VA-specific settings.

Many of the configuration settings for the Services Director VA can be updated from the Services Director VA **System** menu.

Updating General Settings

You can change a variety of general settings for Services Director VA from the **System > General Settings** page. Defaults are applied automatically when the Services Director VA is created. You only need to update these settings to fine-tune the Services Director VA to your specific requirements.

Apply any changes to put them into use immediately.

FIGURE 75 General Settings Page

General Settings

Monitoring

Controller Failure Period:	<input type="text" value="180"/>	Instance Failure Period:	<input type="text" value="180"/>
Controller Monitor Interval:	<input type="text" value="60"/>	Instance Monitor Interval:	<input type="text" value="60"/>
Host Failure Period:	<input type="text" value="180"/>	Monitor Email Interval:	<input type="text" value="60"/>
Host Monitor Interval:	<input type="text" value="60"/>	Overdue Warning Period:	<input type="text" value="300"/>

Metering

Meter interval:

Log check interval:

SNMP enabled: ☒

Licensing

Alert threshold:

Alert threshold interval:

Logging

License logging:

Metering logging:

Inventory logging:

Authentication logging:

Metering logging:

Inventory logging:

Authentication logging:

Monitoring logging:

Backup logging:

Deployment

Max instances:

Bandwidth Licensing

Expire Warning Days:

Controller Licensing

Expire Warning Days:

Instance Registration

Time Out Period: **Hours**

Validate Owners: ☒

Telemetry

Services Director can collect and export anonymized usage data to Pulse Secure. See [this Knowledge Base article](#) for more information.

Enabled: ☒

Apply

Revert

Flexible Licensing Check

FLA Check Status: **Enabled**

Enable

Disable

Metering Alerts and Notifications

Metering Alerts and Notifications Status: **Enabled**

Enable

Disable

Auto Cleanup vTMs

Auto Cleanup vTMs Status: **All vTMs**

Off **Self Registered Auto Accepted** **All**

Services Director can automatically mark instances as deleted if they repeatedly fail monitoring

Updating Monitoring Settings

The following settings enable you to configure monitoring.

- **Controller Failure Period** - the period of time, in seconds, after which a Services Director is considered to have failed. The default value is 180.
- **Controller Monitor Interval** - the period of time, in seconds, between monitoring the Services Director. The default value is 60.
- **Host Failure Period** - the period of time, in seconds, after which a host is considered to have failed. The default value is 180.
- **Host Monitor Interval** - the period of time, in seconds, between monitoring hosts. The default value is 60.
- **Instance Failure Period** - the period of time, in seconds, after which the instance is considered to have failed. The default value is 180.

Note: This period is also used by the automatic deletion of self-registered vTMs, see [“Configuring Auto Cleanup of Virtual Traffic Managers” on page 204](#).

- **Instance Monitor Interval** - the length of the *monitoring cycle*. That is, the period of time, in seconds, between each Services Director attempt to retrieve monitoring information from each vTM. The default value is 60.

Note: This interval is also used by the automatic deletion of self-registered vTMs, see [“Configuring Auto Cleanup of Virtual Traffic Managers” on page 204](#).

- **Monitor Email Interval** - the period of time, in seconds, between monitoring alert emails. The default value is 60.
- **Overdue Warning Period** - the period of time, in seconds, to consider monitoring overdue. The default value is 300.

Updating Metering Settings

The following settings enable you to configure metering.

- **Meter Interval** - the period of time, in seconds, between metering actions. The range is from 1-3600. The default value is 3600 seconds (1 hour).
- **Log Check Interval** - the period of time, in seconds, between checks for log space. The range is from 1-3600. The default value is 3600 seconds (1 hour).
- **SNMP enabled** - this check box is used to enable/disable the use of SNMP. SNMP is used to gather certain types of information (such as metering) from the Virtual Traffic Managers (vTMs) in the estate of the Services Director.

Note: You can monitor the capacity of the `/data` partition that is used for the storage of metering logs. See [“Monitoring the Storage Capacity of Metering Logs” on page 433](#).

Updating Licensing Settings

The following settings enable you to configure licensing.

- **Alert Threshold** - the number of alerts that sent. The range is from 1-3600. The default is 1.
- **Alert Threshold Interval** - the period of time, in seconds, between alerts. The range is from 1-3600. The default value is 3600 seconds (1 hour).

The threshold and interval settings enable you to determine how many requests have to be received by a non-primary license server in the specified interval before an alert email is sent. After the threshold and interval is reached, an alert message is sent. At most, one message is sent per hour, to protect against a flood of messages being sent in the case of complete failure of the primary license server on a busy system.

Updating Logging Settings

The following settings enable you to configure logging.

- **License Logging** - a license value. The range is from 0-10.
 - The default value is 0, which equals no logging.
 - A log level of 3 or higher causes responses to license server requests to be logged in full, including the feature values set by the feature pack and bandwidth associated with the instance making the request.
- **Metering Logging** - the metering logging value. The range is from 0-10.
 - The default value is 0, which equals no logging.
 - A log level of 5 or higher gives a summary of the activities of the metering thread (that is, starting metering, stopping metering, and so forth)
 - A log level of 9 or higher provides a detailed logging of each instance being metered.

Note: You can monitor the capacity of the `/data` partition that is used for the storage of metering logs. See [“Monitoring the Storage Capacity of Metering Logs” on page 433](#).

- **Inventory Logging** - the metering logging value. The range is 0-10.
 - The default value is 0, which equals no logging.
 - A log level of 1 or higher will cause inventory changes to be logged (the equivalent of the audit records).
 - A log level of 3 or higher causes logging of all deployment and action commands.
 - A log level of 8 or higher causes logging of the output from all deployment and actions.

Updating Deployment Settings

The following setting enables you to configure deployment.

- **Max Instances** - the maximum number of vTM instances that can be deployed. The default value is 0, which equals no limit. Typically, this is the correct value for most deployments. Note that:
 - Instances that have been deleted do not count towards the limit.
 - Instances that have been deployed but are not active (that is, have not been started) do count towards the limit.
 - If you create a new instance in excess of this number, the instance is rejected with an error message.
 - If this property is set to a lower number than the number of currently deployed instances then there is no immediate effect but subsequent deployment requests are rejected.

Updating Bandwidth Licensing Settings

The following setting enables you to configure bandwidth licensing.

- **Expire Warning Days** - the number of days to warn you before the bandwidth license expires. The default value is 30.

Updating Controller Licensing Settings

The following setting enables you to configure controller licensing.

- **Expire Warning Days** - the number of days to warn you before the controller license expires. The default value is 30.

Updating Instance Registration Settings

The following settings enables you to configure self-registration.

- **Time Out Period** - the number of hours before a *Warning* self-registration request will transition automatically to *Blacklisted*. The default is 24.
- **Validate Owners** - enables/disables the mandatory validation of the Owner property during the automatic self-registration of vTMs.

Updating Telemetry Settings

Services Director can collect and export usage data to Pulse Secure. The initial setting for this feature is chosen during the Setup Wizard. However, this setting can be changed at any time.

- To enable this feature, enable the **Enable** check box.
- To disable this feature, clear the **Enable** check box.

Updating Metering Alerts and Notifications Settings

The following setting enables you to configure the reporting of metering issues.

- **Metering Alerts and Notifications** - enables/disables the reporting of metering alerts and notifications. See [“Processing Virtual Traffic Manager Metering Discrepancy Warnings” on page 224](#).

Note: You can monitor the capacity of the `/data` partition that is used for the storage of metering logs. See [“Monitoring the Storage Capacity of Metering Logs” on page 433](#).

Configuring the FLA Checker

The Services Director VA uses an automatic FLA checker. Refer to the *Pulse Services Director Advanced User Guide* for details. To configure the global Flexible Licensing Check, click **Enable** or **Disable**. This selection is applied automatically.

Updating Auto Cleanup of Failed vTMs

The following setting enables you to configure the auto-deletion of failed vTMs:

- Auto Cleanup vTMs - enables you to set the required behaviour for the deletion of failed vTMs, see [“Configuring Auto Cleanup of Virtual Traffic Managers” on page 204](#).

Updating Date and Time Settings

You can change the date and time settings for the Services Director VA from the **System > Date and Time Settings** page. Settings are in three categories:

- Basic date and time settings. To change the basic settings, set the correct **Date** and **Time**, and click **Apply**.
- Time zone settings. To change the **Time Zone** for your Services Director, select the required time zone and click **Apply**.
- NTP settings. Where NTP is active, basic date and time settings are overwritten.
 - A default set of NTP services are listed. You can enable or disable any listed service by expanding the service entry and changing its state.
 - You can add another NTP service by clicking **Add** and specifying details for the service.
 - To stop the use of the NTP service, click **Stop**. Click **Start** to restart it.

Updating Administration Credentials

You can change the administration credentials for the Services Director VA from the **System > User Credentials** page. These credentials are used as follows:

- To log in to the Services Director VA.
- To access a terminal session for the Services Director, such as when you wish to use the command-line user interface.
- For REST API authentication.

On the **Services Director Credentials** page, specify a **Password** and a password **Confirm** before clicking **Update**. You are required to authenticate using the new credentials.

Updating Email Settings

You can change the email settings for the Services Director VA from the **System > Email Alerts** page. This page enables you to enter email notification details for your Services Director, to ensure that you receive email notifications for events and failures. You must specify:

- **SMTP Server** - This is either the hostname or IP address of the SMTP server in your network.
- **SMTP Port** - Typically, you will use the default port number, 25.
- **Notification Email** - All email from the Services Director will go to each entry in this comma-separated list of e-mail addresses.
- **From Address** - The required "from" address for all emails.

You can use "\$fqdn" to substitute in this appliance's fully-qualified domain name.

Note: Services Director VA automatically restarts the Services Director service after email changes are applied.

Updating the SSL Certificate

You can replace the SSL certificate for the Services Director VA from the **System > Service SSL Certificate** page. Under **Certificate installed**, click the hyperlink, and select one of the following options:

- **Single file with public and private keys**. Then, click **Choose File** to locate the certificate file.
- **Separate public and private key files**. Then, click **Choose File** to locate each file.
- **Text content of the public and private keys**. Then, paste the required text in.

Apply these changes to put them into use immediately.

Updating the REST API Port

You can update the REST API port used by the Services Director VA from the **System > Service Status** page. Apply this change to put the new port number into use immediately.

You can also start, stop and restart the Services Director service from this page. See [“Starting and Stopping the Services Director Service” on page 418](#).

Updating Security Settings

You can change the security settings for Services Director VA from the **System > Security Settings** page. Defaults are applied automatically when the Services Director VA is created.

This page supports the following functions:

- Changing the Master Password for your Services Director. See [“Changing the Master Password for the Services Director VA” on page 121](#).
- Enabling shell access for command line users of the Services Director. Refer to the *Pulse Services Director Advanced User Guide*.

You can also define the suspension criteria for failed Services Director logins.

FIGURE 76 Security Settings Page: Suspension Settings

Login Settings



Max login attempts:

User lockout duration: Minutes

The **Max login attempts** defines the maximum number of failed Services Director login attempts for a user. Zero (the default setting) indicates that there is no maximum.

If the **Max login attempts** limit is reached, a lockout defined by the **User lockout duration** is applied. This has a default of 1 minute, and a maximum of 1440 minutes (equal to one day).

After the lockout period has ended, the same user can continue to attempt to log in.

Changing the Master Password for the Services Director VA

The master password for the *Active Services Director* VA can be changed from the **Security Settings** page.

Note: If you wish to reset the master password (that is, you do not know what the current master password is), refer to the *Pulse Services Director Advanced User Guide*.

Changing the Master Password

The master password for the *Active Services Director* VA can be changed from the **Security Settings** page.

Note: If you wish to reset the master password (that is, you do not know what the current master password is), refer to the *Pulse Services Director Advanced User Guide*.

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **System** Menu, then click **Security**. The **Security Settings** page appears.

FIGURE 77 Security Settings Page

Master Password

Brocade Services Director uses a master password to encrypt sensitive data. The master password is already set. If you would like to change the password, please enter the details below.

Current Password	<input type="password"/>	Generate Password
New Password	<input type="password"/>	
Confirm Password	<input type="password"/>	

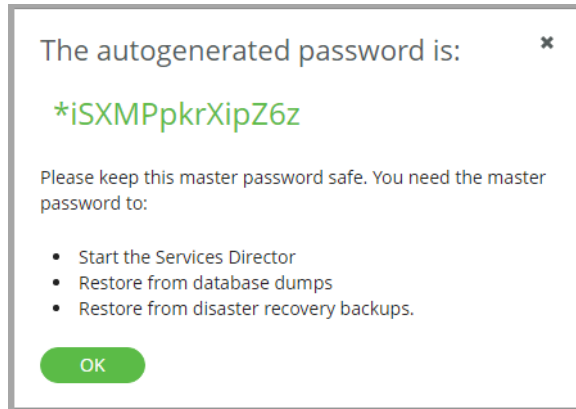
For security, it is recommended that this password is input manually every time the Services Director starts. However, the password could be stored in a file (which is a less secure option but allows for non-interactive start up).

☐ Store the password to a file.

4. Enter the **Current Password**.

5. To change the master password, perform one of the following operations.
 - Enter a new password and confirm the password.
 - Click **Generate Password**. The **Password** and **Confirm Password** fields are populated automatically and an information dialog box is displayed.

FIGURE 78 Autogenerated Password Dialog Box



6. Click **OK** to close the information dialog box after recording the password, and then confirm that you have stored the password in the next dialog box.

It is essential that the master password (whether chosen yourself or generated automatically) is recorded and can be retrieved. Pulse Secure recommends that this password is recorded in a secure location that is separate from the Services Director VA.

7. Choose whether to store the password internally for automatic use:
 - Select the **Store the password in a file** check box to store the master password within the Services Director VA. The password will be automatically available whenever the Virtual Machine for a Services Director VA restarts. However, you must enter the master password manually when you create a Services Director VA from a backup file.
 - Clear the **Store the password in a file** check box to not record the master password. You must to enter the master password manually whenever the Virtual Machine for a Services Director VA restarts, and when you create a Services Director VA from a backup file.

See **“Entering the Master Password After a Virtual Machine Restart” on page 418** for details of restarting a VM.

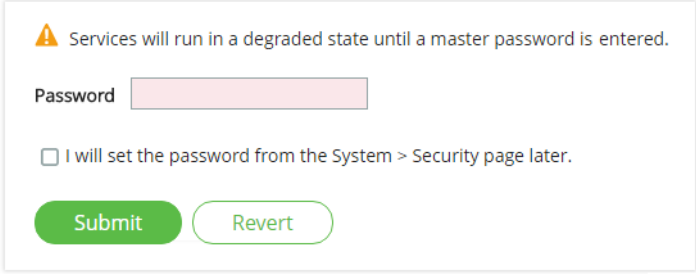
8. Select the **Store the password to a file** check box if you want to store the master password internally for future use.

If you do not choose to store this password, you must enter it after the Virtual Machine for this Services Director VA restarts (see **“Entering the Master Password After a Virtual Machine Restart” on page 418**).

9. Click **Update**. The master password is changed.
10. Access your *Standby* Services Director VA from a browser.
11. Log in as the administration user.

A dialog box requesting the new master password immediately appears:

FIGURE 79 Master Password Required



The dialog box has a light gray border and a white background. At the top left is a yellow warning triangle icon. To its right is the text: "Services will run in a degraded state until a master password is entered." Below this is a label "Password" followed by a red rectangular input field. Underneath the input field is a checkbox with the text "I will set the password from the System > Security page later." At the bottom are two buttons: a green "Submit" button and a white "Revert" button with a green border.

You may receive an e-mail notification of a raised `master_password_fail` alarm between you changing the master password on the *Active* Services Director VA and entering the new master password on the *Standby* Services Director VA.

12. Enter the new master password and click **Submit**.

Adding Virtual Traffic Managers to the Services Director

- [Overview: Adding Virtual Traffic Managers to the Services Director](#) 125
- [Working with vTM Communications Channel](#) 126
- [Adding Resources Required for Virtual Traffic Managers](#) 128
- [Registering an Externally-Deployed Virtual Traffic Manager](#)..... 153
- [Self-Registering an Externally-Deployed Virtual Traffic Manager](#)..... 169
- [Self-Registering a Cloud-Based Virtual Traffic Manager](#)..... 186

Overview: Adding Virtual Traffic Managers to the Services Director

The Services Director supports several methods for adding a Virtual Traffic Manager (vTM) to the estate of the Services Director:

- By registering an externally-deployed vTM from the Services Director. See [“Registering an Externally-Deployed Virtual Traffic Manager” on page 153](#).

Note: This method is not supported for vTMs that use vTM Communications Channel, see [“Working with vTM Communications Channel” on page 126](#).

- By processing a self-registration request that was received from an externally-deployed vTM by the Services Director. See [“Self-Registering an Externally-Deployed Virtual Traffic Manager” on page 169](#).

Note: This method is required for all vTMs that use vTM Communications Channel, including those that are behind a NAT device, see [“Working with vTM Communications Channel” on page 126](#).

- By deploying a vTM from the Services Director VA using an instance host. See the *Pulse Services Director Advanced User Guide* for full details.

Before you perform any of these methods, you must create any required resources, see [“Adding Resources Required for Virtual Traffic Managers” on page 128](#).

The communication between the vTM and the Services Director depends on whether vTM Communications Channel is enabled, see [“Working with vTM Communications Channel” on page 126](#).

Working with vTM Communications Channel

The method of communication between the vTM and the Services Director depends on whether vTM Communications Channel (Comms Channel) is enabled.

Comms Channel is an update of the (pre-19.1) mechanism that enabled communication between each vTM and the Services Director. Comms Channel is only supported on vTMs at v19.1 or later.

The use of Comms Channel only affects the communication between the vTM and the Services Director. When Comms Channel is enabled on a vTM:

- The vTM and the Services Director always use a mutually-authenticated, TLS-based link initiated by the vTM.
- The vTM can be located in a private network behind a NAT device, see [“Enabling a vTM Cluster To Operate Behind a NAT Device” on page 126](#).
- The vTM will always communicate with the *Active* node of an HA pair only.
- The vTM must be self-registered, see [“Self-Registering an Externally-Deployed Virtual Traffic Manager” on page 169](#).

Use of the Comms Channel is the default for self-registered vTMs of 19.1 or later. However, Comms Channel can be disabled if required, see [“Disabling Comms Channel on a vTM” on page 127](#).

Note: The Comms Channel configuration of a vTM is not replicated to all vTMs in a cluster.

Enabling a vTM Cluster To Operate Behind a NAT Device

For vTMs running v19.1 (and later), vTMs may be located in a private network behind a NAT device.

To set up a vTM cluster behind a NAT device:

- All vTMs in the cluster must have Comms Channel enabled, see [“Working with vTM Communications Channel” on page 126](#).
- The vTM cluster must be formed on each vTM using its user interface.
- Each vTM in the cluster must be added to the estate of the Services Director using self-registration from the vTM user interface. Both manual and automatic self-registration methods are supported. See [“Self-Registering an Externally-Deployed Virtual Traffic Manager” on page 169](#).

Disabling Comms Channel on a vTM

When you perform the vTM configuration wizard, if the vTM is at v19.1 or later, it will automatically be configured to use Comms Channel. If required, you can later change the configuration of the vTM so that it does not use Comms Channel. This requires you to re-register the vTM.

If you want to disable Comms Channel:

1. Log into the vTM.
2. Go to **System > Licenses > Services Director Registration**.
3. Set **remote_licensing!comm_channel_enabled** to *NO*.
4. Enable the **Force re-registration** check box.
5. Click **Save and Register**.

The vTM will reconfigure to disable Comms Channel, and re-connect to the Services Director in that mode.

Note: The Comms Channel configuration of a vTM is not replicated to all vTMs in a cluster.

You can enable Comms Channel at any point by repeating this process, and setting **remote_licensing!comm_channel_enabled** to *YES*.

Adding Resources Required for Virtual Traffic Managers

Before you attempt to register any vTM, you must ensure that all required resources are present on the Services Director. The tasks required will vary according to your specific configuration.

- Add any additional licenses. For example, a Resource License to support vTM analytics or additional bandwidth. See [“Adding a License to the Services Director” on page 128](#).
- Create any required Feature Packs, see [“Adding a Feature Pack to the Services Director” on page 129](#).
- Create any required Owner entries, see [“Adding an Owner to the Services Director” on page 141](#).
- Create any required Legacy licenses, see [“Adding a Legacy FLA License to the Services Director” on page 144](#).
- Create any required Access Profiles, see [“Creating an Access Profile \(vTM User Authentication Only\)” on page 272](#).

Adding a License to the Services Director

The functionality of the Services Director is determined by three kinds of licenses, and the Stock Keeping Units (SKUs) identified by these licenses:

- The Services Director License. This major license enables the use of the Services Director.

The SKU identified by this license defines the customer type (Enterprise or CSP), the Feature Tier and the individual functions that are available in the Services Director. The SKU is central to the creation of a Feature Pack for use on external vTMs.
- Resource Licenses. These secondary licenses enable the use of limited resources on the Services Director by an Enterprise customer.

The SKU identified by a Resource License is typically for Bandwidth allocation or vTM Analytics features, and is added to a Feature Pack to make the resource available to any vTM that uses the Feature Pack.
- Add-on Licenses. These are historical licenses associated with “old style” Services Director licenses. They were used on the Services Director by Enterprise customers only.

Note: Add-On Licenses are incompatible with “new style” Services Director licenses.

Note: Universal FLA Licensing and Legacy FLA Licensing are also supported, but these are used by the vTMs for licensing purposes only. See [“Adding a Legacy FLA License to the Services Director” on page 144](#).

Note: To create a Feature Pack, see [“Adding a Feature Pack to the Services Director” on page 129](#).

You add and view licenses from the **Licenses** page.

FIGURE 80 The Licenses Page

Licenses

Services Director Licenses

+ Add

License Key ↕	Valid From ↕	Valid Until ↕	Status ↕
▶ LK1-BR_ADC_MGMT_STDBASE_S_01:857105-0000-43FD-5-0D21-926A-4186	Perpetual	2017-08-27	Active

Resource Licenses

+ Add

License Key ↕	Valid From ↕	Valid Until ↕	Status ↕	SKU ↕
▶ LK1-BR_ADC_FLEX_ADV5G_S_01:1388:377966:20170817T2108011503029281-0000-43FD-5-5CDD-621E-D477	Perpetual	2017-08-27	Active	ENT-ADVANCED
▶ LK1-BR_ADC_RES_EMBASSI_S_01:5:186105:20170817T2108011503029281-0000-43FD-5-9A73-DD73-33CE	Perpetual	2017-08-27		ENT-ENTM

Add-on Licenses

+ Add

Add-on License Key ↕	Valid From ↕	Valid Until ↕
No Data		

The process for adding additional licenses is similar for all license types:

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The Home page appears.
3. Click the **Catalogs** menu, and then click **Licensing: Licenses**. The **Licenses** page appears.
4. Click the **Add** plus symbol for your required license type. A licensing dialog box window appears. For example, for Resource Licenses:

FIGURE 81 The Licenses Page: Adding a License

5. Enter the license number and click **Add**.

The new license is added in its category in the **Licenses** page.

After all new licenses are added, create one or more Feature Packs that include them. See [“Adding a Feature Pack to the Services Director” on page 129](#).

Note: Existing Feature Packs cannot be updated.

Adding a Feature Pack to the Services Director


Before you register any vTM instances, you must define one or more Feature Packs.

A Feature Pack defines the Services Director features that are available to a vTM instance once you have registered it on the Services Director.

The total set of features that are available in a Feature Pack is defined by its selected *Feature Tier*.

- Each Feature Tier is a subset of the tier above it.
- Feature Tiers include features that are relevant to your license type: Enterprise or Cloud Service Provider (CSP).
- Enterprise licenses have access to *Advanced* and *Enterprise* tiers only.
- CSP licenses have access to *Basic*, *Standard*, *Advanced* and *Enterprise* tiers.

FIGURE 82 Features Tiers for Enterprise and CSP licenses



Note: The *Enterprise* feature tier should not be confused with the Enterprise customers/licenses, or Analytics Resource Pack Licenses.

For CSP licenses only, a Feature Pack also requires:

- A bandwidth, expressed as either Mbps or Gbps.
- A pricing model - *Fixed Price Monthly*, *Fixed Price Weekly*, or *Hourly plus Data Transfer*.

Once all Feature Pack properties are defined, the system is able to identify the Stock-Keeping Unit (SKU) that is required for the Feature Pack. You can exclude any of the SKU's features from the Feature Pack if required.

Enterprise customers can include extra SKUs from one or more Resource Licenses to augment the base SKU. For example, to add vTM Analytics features. See [“Adding a License to the Services Director” on page 128](#).

Note: A list of features for a SKU can be seen on the expanded view of a SKU in the **SKUS and Feature Packs** page.

A default Feature Pack (typically a SKU with no exclusions) is created automatically when you install the Services Director VA based on an Enterprise license.

The procedure for creating a Feature Pack is dependent on your license type.

- For current Enterprise licenses, see [“Adding a Feature Pack for an Enterprise License” on page 134](#).
- For current Cloud Service Provider (CSP) licenses, see [“Adding a Feature Pack for a CSP License” on page 131](#).
- For older Enterprise/CSP licenses, see [“Adding a Feature Pack for an Older License” on page 136](#).

Adding a Feature Pack for a CSP License

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **SKUS and Feature Packs**.

The **SKUS and Feature Packs** page appears.

FIGURE 83 SKUS and Feature Packs Page: CSP

SKUs and Feature Packs

Feature Packs

[+ Add](#)

Feature Pack Name	SKU	Add-on SKUs	Status	Info	Actions
▶ CSP_full	BR-ADC-UTLM-ADV10M-U-01		Active	No exclusions	Apply

SKUs

Show only compatible SKUs ☒

SKU Name	Details	Compatible	Status
▶ BR-ADC-UTLH-ADV10M-U-01	CSP Advanced Hourly 10Mbps	✓	Active
▶ BR-ADC-UTLH-ADV1G-U-01	CSP Advanced Hourly 1Gbps	✓	Active
▶ BR-ADC-UTLH-ADV300M-U-01	CSP Advanced Hourly 300Mbps	✓	Active

4. Select the **Show only compatible SKUs** check box to ensure that only SKUs that are compatible with your license are displayed.
5. In the table of SKUs, locate the SKU from which you wish to create a Feature Pack.

- Expand this SKU to view its supported features. For example, the *BR-ADC-UTLM-ADV100M-U-01* SKU:

FIGURE 84 SKUS and Feature Packs Page: Expanded SKU

▼	BR-ADC-UTLM-ADV100M-U-01	CSP Advanced Monthly 100Mbps	✓	Active
SKU Name:	BR-ADC-UTLM-ADV100M-U-01			
Details:	CSP Advanced Monthly 100Mbps			
Pricing Model:	fixed price monthly			
Feature Tier:	Advanced			
Fixed Resource Usage:	100 Mbps			
Compatible:	✓			
Status:	Active			
Features:	<div> <div>anlyt</div> <div>Enable Realtime Analytics.</div> </div> <div> <div>auto</div> <div>Enable Autoscaling.</div> </div> <div> <div>bwm</div> <div>Enable Bandwidth Management classes.</div> </div> <div> <div>cache</div> <div>Enable Web Caching.</div> </div>			

- Locate the feature(s) that you wish to exclude, and make a note of the feature name. For example, the *auto* (Autoscaling) feature. That is, this Feature Pack will not support the Autoscaling feature. All other features will still be supported.
- Collapse the SKU in the table.
- Click the **Add** button above the table of feature packs.

The **Add Feature Pack** dialog box appears.

FIGURE 85 SKUS and Feature Packs Page: Add Feature Pack

Add Feature Pack

Feature Pack Name:

Pricing Model:

☐ Fixed Price Monthly
 ☐ Fixed Price Hourly

Feature Tier:

SKU Code:

Excluded:

Add-on SKUs:

N/A

Info:

Add

- Enter a **Feature Pack Name**.

This name will appear in the table of Feature Packs.

11. Select a **Pricing Model**.
12. Select the required **Feature Tier**.
13. Select a **Bandwidth**.

The displayed SKU Code updates automatically to reflect your choices.

14. Enter a space-separated list of **Excluded** features.
15. Enter a description for the Feature pack as **Info**.

This name will appear in the table of Feature Packs.

FIGURE 86 SKUS and Feature Packs Page: Specify New Feature Pack

Add Feature Pack ✕

Feature Pack Name:

Pricing Model: ☒ Fixed Price Monthly ☐ Fixed Price Hourly

Feature Tier:

Bandwidth:

SKU Code: BR-ADC-UTLM-ADV10M-U-01

Excluded:

Add-on SKUs: N/A

Info:

16. Click **Add**. The new Feature Pack is added to the table of Feature Packs.

FIGURE 87 SKUS and Feature Packs Page: New Feature Pack Added

Feature Packs

Feature Pack Name ↕	SKU ↕	Add-on SKUs ↕	Status ↕	Info ↕	Actions
▶ CSP_full	BR-ADC-UTLM-ADV10M-U-01		Active	No exclusions	<input type="button" value="Apply"/>
▶ CSP_not_auto	BR-ADC-UTLM-ADV10M-U-01		Active	No autoscaling	<input type="button" value="Apply"/>

17. (Optional) Expand the Feature Pack to see its full details.
18. (Optional) You can apply this new Feature Pack to one or more registered instances, see [“Applying a Feature Pack to Registered Instances” on page 140](#).
19. Repeat this process to create all required Feature Packs.

Adding a Feature Pack for an Enterprise License

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **SKUS and Feature Packs**.

The **SKUS and Feature Packs** page appears.

FIGURE 88 SKUS and Feature Packs Page: Enterprise
SKUs and Feature Packs

Feature Packs

 Add

Feature Pack Name ▾	SKU ▾	Add-on SKUs ▾	Status ▾	Info ▾	Actions
▶ ENT-ADVANCED_full	ENT-ADVANCED		Active	No exclusions	<button>Apply</button>

SKUs

Show only compatible SKUs ☒

SKU Name ▾	Details ▾	Compatible	Status ▾
▶ ENT-ADE	Data Export	✓	Active
▶ ENT-ADVANCED	ENT Advanced	✓	Active
▶ ENT-ANALYTICS	Analytics	✓	Active
▶ ENT-ENTERPRISE	ENT Enterprise	✓	Active
▶ ENT-ENTM	Enterprise Management	✓	Active
▶ ENT-WAFPROXY	ENT WAFProxy	✓	Active

4. Select the **Show only compatible SKUs** check box to ensure that only SKUs that are compatible with your license are displayed.
5. In the table of SKUs, locate the SKU from which you wish to create a Feature Pack.

- Expand this SKU to view its supported features. For example, the ENT-ADVANCED SKU:

FIGURE 89 SKUS and Feature Packs Page: Expanded SKU

▼	ENT-ADVANCED	ENT Advanced	✓	Active
SKU Name:	ENT-ADVANCED			
Details:	ENT Advanced			
Pricing Model:	prepaid			
Feature Tier:	Advanced			
Fixed Resource Usage:	N/A			
Compatible:	✓			
Status:	Active			
Features:	<div> <div>anlyt</div> <div>Enable Realtime Analytics.</div> </div> <div> <div>auto</div> <div>Enable Autoscaling.</div> </div> <div> <div>bwm</div> <div>Enable Bandwidth Management classes.</div> </div> <div> <div>cache</div> <div>Enable Web Caching.</div> </div>			

- Locate the feature(s) that you wish to exclude, and make a note of the feature name. For example, the *cache* (Web Caching) feature. That is, this Feature Pack will not support the Web Caching feature. All other features will still be supported.
- Collapse the SKU in the table.
- Click the **Add** button above the table of feature packs.

The **Add Feature Pack** dialog box appears.

FIGURE 90 SKUS and Feature Packs Page: Add Feature Pack

Add Feature Pack

Feature Pack Name:

Feature Tier:

Advanced

SKU Code:

ENT-ADVANCED

Excluded:

Add-on SKUs:

☐ ENT-ANALYTICS

Info:

Add

- Enter a **Feature Pack Name**.

This name will appear in the table of Feature Packs.

- Select the required **Feature Tier**.
- Enter a space-separated list of **Excluded** features.

- Optionally, select one or more **Add-on SKUs**. Each such SKU adds an additional resource (such as Analytics) to the base **SKU Code**.

In this example, an Analytics Resource Pack license has already been added to the Services Director to enable the use of vTM Analytics (see [“Working with vTM Analytics” on page 279](#)). The *ENT-ANALYTICS* SKU is made available by the Analytics Resource Pack license, and you can add this add-on SKU to the Feature Pack to augment the base SKU with analytics capability.

FIGURE 91 SKUS and Feature Packs Page: Specify New Feature Pack

- Optionally, enter a description for the Feature Pack as **Info**.

This name will appear in the table of Feature Packs.

- Click **Add**. The new Feature Pack is added to the table of Feature Packs.

FIGURE 92 SKUS and Feature Packs Page: New Feature Pack Added

Feature Packs

+ Add

	Feature Pack Name ▾	SKU ▾	Add-on SKUs ▾	Status ▾	Info ▾	Actions
▶	ENT-ADVANCED_full	ENT-ADVANCED		Active		Apply
▶	ENT-ADV-Analytics	ENT-ADVANCED	ENT-ENTM	Active	Includes vTM Analytics	Apply

- (Optional) Expand the Feature Pack to see its full details.
- (Optional) You can apply this new Feature Pack to one or more registered instances, see [“Applying a Feature Pack to Registered Instances” on page 140](#).
- Repeat this process to create all required Feature Packs.

Adding a Feature Pack for an Older License

- Access your Services Director VA from a browser, using its Service Endpoint IP Address.
- Log in as the administration user. The **Home** page appears.
- Click the **Catalogs** menu, and then click **SKUS and Feature Packs**.

The **SKUS and Feature Packs** page appears.

FIGURE 93 SKUS and Feature Packs Page

SKUs and Feature Packs

Feature Packs

+ Add

Feature Pack Name ▾	SKU ▾	Add-on SKUs ▾	Status ▾	Info ▾	Actions
▶ STM-400_full	STM-400		Active		Apply

SKUs

Show only compatible SKUs ☒

SKU Name ▾	Details ▾	Compatible	Status ▾
▶ STM-100		✓	Active
▶ STM-200		✓	Active
▶ STM-300		✓	Active
▶ STM-400		✓	Active
▶ STM-WAFPROXY		✓	Active

4. Select the **Show only compatible SKUs** check box to ensure that only SKUs that are compatible with your license are displayed.
5. In the table of SKUs, locate the SKU from which you wish to create a Feature Pack.
6. Expand this SKU to view its supported features. For example, the STM-400 SKU:

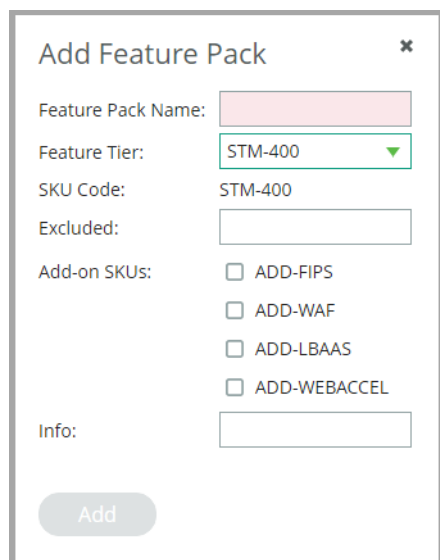
FIGURE 94 SKUS and Feature Packs Page: Expanded SKU

▼ STM-400	✓	Active
SKU Name:	STM-400	
Details:		
Pricing Model:	hourly	
Feature Tier:	STM-400	
Fixed Resource Usage:	N/A	
Compatible:	✓	
Status:	Active	
Features:	<div> anlyt Enable Realtime Analytics. auto Enable Autoscaling. bwm Enable Bandwidth Management classes. cache Enable Web Caching. </div>	

7. Locate the feature(s) that you wish to exclude, and make a note of the feature name. For example, the *Lbrnd* (Random Load Balancing) feature. That is, this Feature Pack will not support the Random load balancing feature. Other load balancing features, such as Round Robin, will still be supported.
8. Collapse the SKU in the table.
9. Click the **Add** button above the table of feature packs.

The **Add Feature Pack** dialog box appears.

FIGURE 95 SKUS and Feature Packs Page: Add Feature Pack



Add Feature Pack ✕

Feature Pack Name:

Feature Tier:

SKU Code:

Excluded:

Add-on SKUs:

- ☐ ADD-FIPS
- ☐ ADD-WAF
- ☐ ADD-LBAAS
- ☐ ADD-WEBACCEL

Info:

Add

10. Enter a **Feature Pack Name**.

This name will appear in the table of Feature Packs.

11. Select the required **Feature Tier**.

This list is defined by the bandwidth packs added to the Services Director.

12. Enter a space-separated list of **Excluded** features.

13. Select any required **Add-on SKUs**.

14. Enter a description for the Feature pack as **Info**.

This description will appear in the table of Feature Packs.

FIGURE 96 SKUS and Feature Packs Page: Specify New Feature Pack

Add Feature Pack

Feature Pack Name:

STM-400_LB

Feature Tier:

STM-400

SKU Code:

STM-400

Excluded:

lbrnd

Add-on SKUs:

☐ ADD-FIPS

☐ ADD-WAF

☐ ADD-LBAAS

☐ ADD-WEBACCEL

Info:

Excl. Random LB

Add

15. Click **Add**. The new Feature Pack is added to the table of Feature Packs.

FIGURE 97 SKUS and Feature Packs Page: New Feature Pack Added

Feature Packs

+

Add

Feature Pack Name	SKU	Add-on SKUs	Status	Info	Actions
▶ STM-400_full	STM-400		Active		<div>Apply</div>
▶ STM-400_LB	STM-400		Active	Excl. Random LB	<div>Apply</div>

16. (Optional) Expand the Feature Pack to see its full details.
17. (Optional) You can apply this new Feature Pack to one or more registered instances, see [“Applying a Feature Pack to Registered Instances” on page 140](#).
18. Repeat this process to create all required Feature Packs.

Once you have created all required Feature Packs, you can use these to register and deploy vTM instances.

Applying a Feature Pack to Registered Instances

Once you have added a Feature Pack, you may want to apply it to one or more registered instances.


1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **SKUS and Feature Packs**.


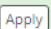
The **SKUS and Feature Packs** page appears. For example:

FIGURE 98 SKUS and Feature Packs Page

SKUs and Feature Packs

Feature Packs

 Add

Feature Pack Name	SKU	Add-on SKUs	Status	Info	Actions
▶ ENT-ADVANCED_full	ENT-ADVANCED		Active		
▶ ADVANCED+EM	ENT-ADVANCED	ENT-ENTM	Active		

4. For the required Feature Pack, click the **Apply** action.

A selection dialog appears. For example:


FIGURE 99 Applying a Feature Pack: Selection of vTMs

Select Instances to apply Feature Pack to

Feature Pack:

☐ Select all


Select	Name	Bandwidth	Cluster	Feature Pack
<input type="checkbox"/>	cerulean-01	50	Cluster-1QFP-Y1AC-UBN3-3SR0	ENT-ADVANCED_full
<input type="checkbox"/>	cerulean-02	50	Cluster-AE75-9ID3-80HY-AU1H	ENT-ADVANCED_full



5. Click the **Select** check box for each vTM to which you want to apply the Feature Pack.
6. Click **Apply**.

7. A completion message appears. For example:

FIGURE 100 Applying a Feature Pack: Complete



8. Close the dialog.
9. (Optional) Confirm the result in the **vTM Instances** page.

FIGURE 101 Confirming the Application of a Feature Pack

vTM Instances

► **Filters** Filtering by Lifecycle, Instance Health, Licensing Health

+ Add Show: 20 ▼ of 2 Instances

Name	License Name	Bandwidth	Feature Pack	Version	Cluster	Instance Lifecycle	Instance Health	Licensing Health	Action
▶ cerulean-01	universal_v4	50	ADVANCED+EM	17.3	Cluster-1QFP-Y1AC-UBN3-3SR0	Active	OK	Licensed	N/A
▶ cerulean-02	universal_v4	50	ADVANCED+EM	17.3	Cluster-AE75-9ID3-80HY-AU1H	Active	OK	Licensed	N/A

« < Page 1/1 > »

Adding an Owner to the Services Director

There are several Services Director resources that require an *owner*. This property identifies a person or organization that is associated with a resource, and optionally includes contact information.

For example, a single owner entry can be used for all resources owned by a Enterprise customer. Alternatively, an owner entry can be created to identify individual customers for resources supplied by a Cloud Service Provider.

The following resources require an owner:


- An externally-deployed vTM instance. See **“Registering an Externally-Deployed Virtual Traffic Manager” on page 153.**
- A vTM instance that is deployed using an instance host. Refer to the *Pulse Services Director Advanced User Guide*.
- A vTM Cluster. See **“Creating a Virtual Traffic Manager Cluster” on page 233.**

Creating an Owner

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **Owners**. The **Owners** page appears.

FIGURE 102 Owners Page

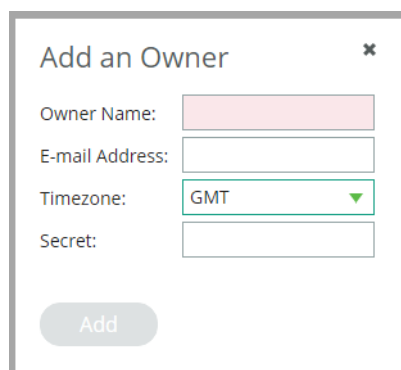
Owners

 Add

Name ▾	Owner ID ▾	E-mail address ▾	Timezone ▾
▶ JK	Owner-KEK0-7VEV-VWD4-YBOM	admin@tk.com	Europe/London
▶ JDDJ	Owner-9SZQ-L514-8KBY-DVLK	admin@judodojo.com	Africa/Asmera
▶ Venkman	Owner-WHK5-VM7B-ZV8B-JOED	admin@firehouse.com	America/New York

4. Click the **Add** button above the table of Owners. The **Add an Owner** dialog appears.

FIGURE 103 Owners Page: Adding an Owner



The dialog box titled "Add an Owner" contains the following fields:

- Owner Name:** A text input field with a red border.
- E-mail Address:** A text input field.
- Timezone:** A dropdown menu currently showing "GMT".
- Secret:** A text input field.
- Add:** A button at the bottom left.

5. Enter an **Owner Name** for the new entry.
6. (Optional) Enter an **E-mail Address** for the owner.
7. Select the required timezone for the owner.
8. (Optional) Enter a **Secret** password for the owner. This is used during self-registration.
9. Click **Add**. The new Owner is added to the table of Owners.
10. Expand an Owner to view its full details, see ["Viewing Full Details for an Owner" on page 143](#).
11. Repeat this process to create all required Owners.


Once you have created all required Owners, you can use these to register and deploy vTMs and vTM clusters.

Viewing Full Details for an Owner

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **Owners**. The **Owners** page appears.
4. Locate and expand an Owner to view its full details. For example:

FIGURE 104 Owners Page: Displaying Full Details for an Owner

Owners


 Add

Name	Owner ID	E-mail address	Timezone
▼ JK	Owner-WUPO-RLBZ-SAPQ-RAM3	jk@demo.com	GMT

Owner Name:

E-mail Address:

Timezone: ▼

Secret: 

Instances: cerulean-01, cerulean-02

Clusters: Cerulean-Cluster

▶ TK	Owner-07RO-HRCL-4Z1K-YWG2	tk@demo.com	UTC
------	---------------------------	-------------	-----

The properties of the Owner are as follows:

- **Owner Name:** The name of the Owner.
- **E-mail Address:** (Optional) The e-mail address for a point of contact (typically, the admin user) for the Owner.
- **Timezone:** The selected timezone for the Owner.
- **Secret:** (Optional) The password for the Owner. This is used during self-registration.
- **Instances:** A list of vTM instances that are associated with the Owner. This is empty if the Owner is not in use.
- **Clusters:** A list of vTM clusters that are associated with the Owner. This is empty if the Owner is not in use.

5. (Optional) Change the Owner's properties and click **Apply** to update the Owner.

Adding a Legacy FLA License to the Services Director

The Pulse Secure Services Director comes with a pre-installed *Universal FLA License*. This is suitable for any vTM at version 10.1 or later with an active REST API. In all other cases, a *Legacy FLA License* is required. That is:

- The vTM version is 10.0 or earlier.
- The vTM (any version) has its REST API disabled.


You can install a Legacy FLA License using the Services Director VA, after which you can install either of these vTM types. This procedure can also be used to update a Legacy FLA license to a Universal FLA License.

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **Licensing: FLA Licenses**. The **Flexible Licenses** page appears.

When the Services Director is first installed, only the pre-installed Universal FLA License is shown on this page; no Legacy FLA Licenses are present.

FIGURE 105 Flexible Licenses Page: No Legacy FLA

FLA Licenses

 Add License

Universal Licenses

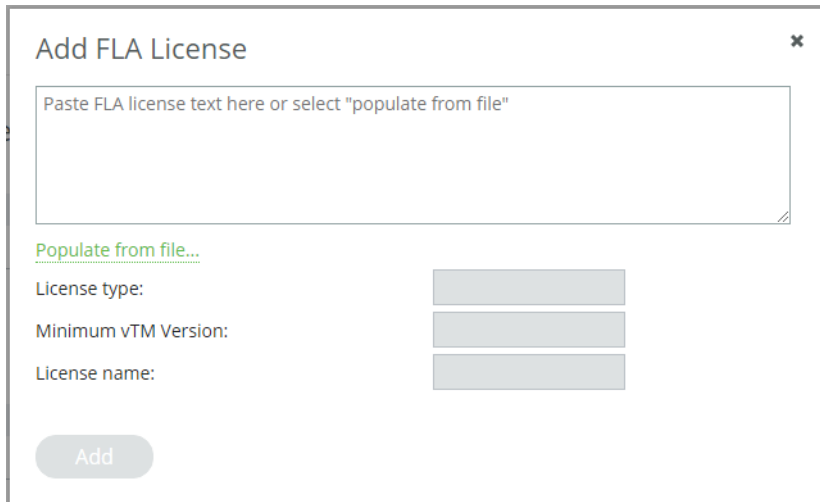
License Name ▾	Status ▾	Default ▾	Actions
▶ universal_v4	Active	Yes	<button>Relicense</button>

Legacy Licenses

License Name ▾	Status ▾	Default ▾	Actions
No Data			

- Click the **Add License** plus symbol. A licensing dialog box window appears.

FIGURE 106 Add FLA License Dialog Box

A dialog box titled "Add FLA License" with a close button (X) in the top right corner. It contains a large text area with the placeholder text "Paste FLA license text here or select 'populate from file'". Below the text area is a link "Populate from file...". Underneath are three input fields: "License type:", "Minimum vTM Version:", and "License name:". At the bottom left is a button labeled "Add".

Add FLA License

Paste FLA license text here or select "populate from file"

[Populate from file...](#)

License type:

Minimum vTM Version:

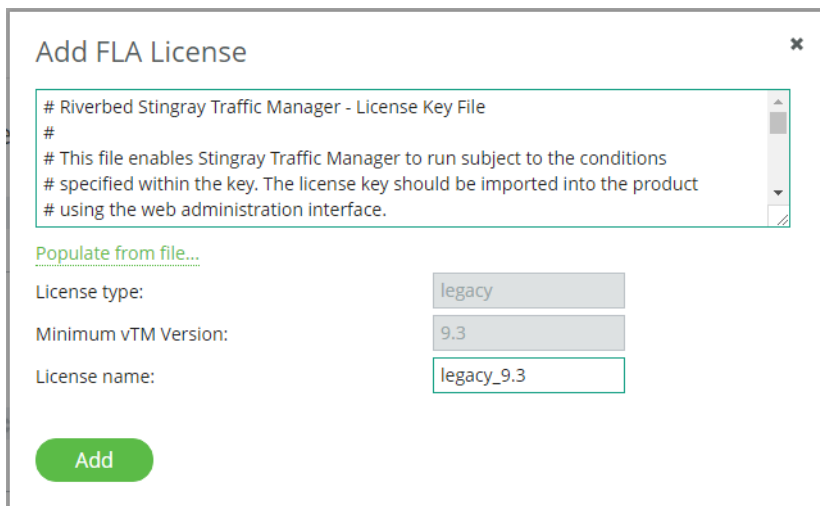
License name:

Add

- Then, either:
 - Paste the text of the Legacy FLA License into the text box, OR
 - Click **Populate from File**, select the file and then click **Upload**. This will populate the text box.

The remainder of the fields in the dialog box will then update to provide license information:

FIGURE 107 License Information

The same "Add FLA License" dialog box as in Figure 106, but now the text area is populated with license information. The "License type" field is set to "legacy", "Minimum vTM Version" is "9.3", and "License name" is "legacy_9.3". The "Add" button is now green.

Add FLA License

Riverbed Stingray Traffic Manager - License Key File

This file enables Stingray Traffic Manager to run subject to the conditions
specified within the key. The license key should be imported into the product
using the web administration interface.

[Populate from file...](#)

License type:

Minimum vTM Version:

License name:

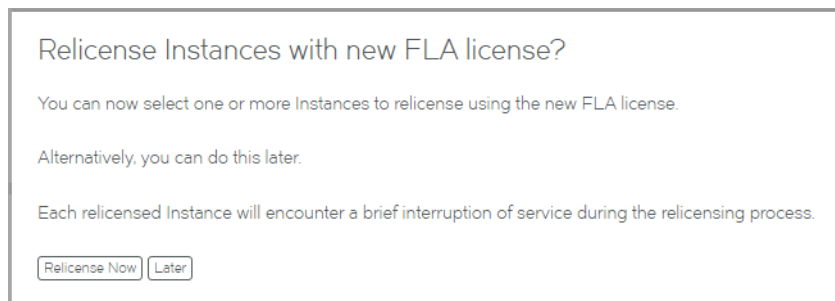
Add

6. Click **Add**.

A relicensing dialog box appears. This enables you to apply the new Legacy FLA License to vTM instances that are currently using a different Legacy FLA License.

See **“Relicensing Virtual Traffic Managers” on page 221** for details of the FLA relicensing mechanism.

FIGURE 108 Relicensing Dialog Box

7. Click **Later**.

You can perform relicensing operations from the **FLA Licenses** page.

The new license is added to the **FLA Licenses** page.

FIGURE 109 Flexible Licenses Page: Legacy FLA Added

FLA Licenses

[+ Add License](#)

Universal Licenses

License Name ▾	Status ▾	Default ▾	Actions
▶ universal_v4	Active	Yes	Relicense

Legacy Licenses

License Name ▾	Status ▾	Default ▾	Actions
▶ legacy_9.3	Active	Yes	Relicense

- Repeat this procedure if you require additional licenses.
- Both Legacy FLA Licenses and Universal FLA Licenses have a default FLA. If you have more than one FLA license for either type, and want to make it the default license for that type, click **Make Default**.

Adding an Auto-Accept Policy to the Services Director

If you want to configure vTMs for automatic self-registration, you will need to create one or more auto-accept policies.

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **Auto-Accept Policies**. The **Auto-Accept Policies** page appears.

FIGURE 110 Auto-Accept Policies Page

Auto-Accept Policies

+ Add								
Name	Policy ID	Management Subnet	Bandwidth (Mbps)	Feature Pack	Accepted Versions	Access Profile	Analytics Profile	
▶ Cerulean	Policy-G5FV-70V2-OLV9-VCYM	255.255.192.0/18	50	ENT-ADVANCED_full	11.1 - 17.3	None	None	

4. Click the **Add** button above the table of auto-accept policies. The **Add an Auto Accept Policy** dialog appears.

FIGURE 111 Auto-Accept Policies Page: Adding an Auto-Accept Policy

Add an Auto-Accept Policy

Policy Name:

Management IP subnet:

Feature Pack:

Bandwidth:

Minimum Version:

Maximum Version:

Access Profile:

Analytics Profile:

Add

5. Enter a unique **Policy Name** for the auto-accept policy.
6. Enter a **Management IP subnet** for the auto-accept policy. This identifies the subnet to which a vTM must belong to be accepted by this policy.

If a vTM that is evaluated by this policy is from outside this subnetwork, the auto-acceptance of the vTM is rejected by the auto-accept policy.

7. Select a **Feature Pack** for the auto-accept policy. This is the feature pack that will be assigned to a vTM that is successfully evaluated using this policy.

This is not an acceptance condition, but the evaluation of the **Bandwidth** property refers to this property.

8. Enter the **Bandwidth** for the auto-accept policy. This is the required bandwidth for a vTM that is evaluated using this policy.

If there is insufficient bandwidth in the specified **Feature Pack** for a vTM, the auto-acceptance of the vTM is rejected by the auto-accept policy.

9. (Optional) Select a **Minimum Version** for the vTM software. This takes the form X.Y. Examples: 10.0, 10.3.

R1 releases are included automatically for any base version. For example, 10.0 includes 10.0r1.

If a vTM that is evaluated by this policy does not meet this condition, the auto-acceptance of the vTM is rejected by the auto-accept policy.

Where a **Minimum Version** is not specified for a policy, the version will be displayed as "Any" in the **Accepted Versions** property in the table of policies.

10. (Optional) Select a **Maximum Version** for the vTM software. This takes the form X.Y. Examples: 10.4, 11.0.

R1 releases are included automatically for any base version. For example, 10.3 includes 10.3r1.

If a vTM that is evaluated by this policy does not meet this condition, the auto-acceptance of the vTM is rejected by the auto-accept policy.

Where a **Maximum Version** is not specified for a policy, the version will be displayed as "Any" in the **Accepted Versions** property in the table of policies.

11. (Optional) Select an **Access Profile**.

This access profile identifies the authenticator and permission groups that will be applied to any vTM that is accepted using this policy.

All cluster members are affected by this change. See ["Working with User Authentication" on page 259](#).

12. (Optional) Select an **Analytics Profile**.

This analytics profile identifies the vTM analytics settings that will be applied to any vTM that is accepted using this policy.

Note: All cluster members are affected by this change. See ["Working with vTM Analytics" on page 279](#).

13. Click **Add**. The new auto-accept policy is added to the table of policies.
14. Expand an auto-accept policy to view its full details.
15. Repeat this process to create all required auto-accept policies.

Once you have created all required auto-accept policies, you can use these to automatically register vTMs, see [“Requesting Self-Registration During vTM Installation” on page 172](#).

Adding a Cloud Registration Resource to the Services Director

If you want to create a cloud-based vTM that will self-register automatically on the Services Director, you must first create a Cloud Registration resource on the Services Director. This process requires you to have AWS login credentials.

Before you create a Cloud Registration resource, you must also create:

- The required Owner on the Services Director, see [“Adding an Owner to the Services Director” on page 141](#).
- The required Auto-Accept Policy on the Services Director, see [“Adding an Auto-Accept Policy to the Services Director” on page 147](#).

Note: You can create a Cloud Registration resource without either an Owner or a Self-Registration Policy property, but the resulting vTM will not contain sufficient information to register automatically on the Services Director. When this happens, you must process the self-registration manually, see [“Processing Self-Registration Requests Manually” on page 182](#).

Once you have created a Cloud Registration resource, you can:

- View the user data text block that is required for the creation of the first cloud-based vTM in a cluster, see [“Viewing User Data Text for a Cloud Registration Resource” on page 151](#).
- Create the first cloud-based vTM in a cluster, see [“Creating a Cloud-Based Virtual Traffic Manager” on page 188](#).

Adding a Cloud Registration Resource

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.

- Click the **Catalogs** menu, and then click **Cloud Registration**. The **Setup Cloud Registration** page appears.

FIGURE 112 Setup Cloud Registration Page

Setup Cloud Registration

AWS Credentials

Set credentials used by Services Director for AWS access. These are required for AWS-based vTMs that will use Traffic IP groups and AWS autoscaling. They can be set now, in which case the generated Cloud Registration User Data will include them, or they may be left blank and then filled in at deployment-time, if required.

If they are left blank then any generated User Data will include placeholders for the keys.

AWS Access Key

AWS Secret Access Key

Update

Revert

AWS Cloud Registrations

+ Add

Name	Owner	Auto-Accept Policy
No Data		

- Specify your AWS credentials as **AWS Access Key** and **AWS Secret Access Key**, and then click **Update**.
- Click the **Add** button above the table of Cloud Registration resources. The **Create a New Cloud Registration** dialog appears.

FIGURE 113 Cloud Registration Page: Adding a Cloud Registration

Create New Cloud Registration

Owner validation is enabled. For auto-acceptance to succeed, a pre-defined Owner must be included in the User Data you provide to AWS. You can disable owner validation on the [Settings](#) page.

Name:

Owner:

Auto-Accept Policy:

Add

- Enter a unique **Name** for the Cloud Registration resource.

- (Optional) Select an **Owner** for the Cloud Registration resource.

If you do not specify an owner before registration, you cannot perform an automatic self-registration of the cloud-based vTM. However, this information can be added in the AWS system before registration.

You can disable the mandatory validation of this property from the **General Settings** page, see [“Updating Instance Registration Settings” on page 117](#).

- (Optional) Select an **Auto-Accept Policy** for the Cloud Registration resource. This is the auto-accept policy that will be used during the evaluation of a cloud-based vTM's self-registration.

If you do not specify an auto-accept policy before registration, you cannot perform an automatic self-registration of the cloud-based vTM. However, this information can be added in the AWS system before registration.

- Click **Add**. The new Cloud Registration resource is added to the table of Cloud Registration resources. For example:

FIGURE 114 Cloud Registration Page: Cloud Registration Added

AWS Cloud Registrations

+ Add		
Name ▾	Owner ▾	Auto-Accept Policy ▾
▶ cloud-reg-01	JK	Accept-Policy-01

- Expand a Cloud Registration resource to view the user data text block that is required for cloud-based registration, see [“Viewing User Data Text for a Cloud Registration Resource” on page 151](#).
- Repeat this process to create all required Cloud Registration resources.

Once you have created a required Cloud Registration resource, you can use it to create the first cloud-based vTM in a cluster, see [“Creating a Cloud-Based Virtual Traffic Manager” on page 188](#).

Viewing User Data Text for a Cloud Registration Resource

The Cloud Registrations page enables you to view and copy the user data text block for individual Cloud Registration resources. This text is required when creating a cloud-based vTM, see [“Creating a Cloud-Based Virtual Traffic Manager” on page 188](#).

- Access your Services Director VA from a browser, using its Service Endpoint IP Address.
- Log in as the administration user. The **Home** page appears.
- Click the **Catalogs** menu, and then click **Cloud Registration**. The **Setup Cloud Registration** page appears.
- In the table of AWS Cloud Registrations, locate the required Cloud Registration entry.

- Expand the Cloud Registration entry to view the user data text block. By default, this uses base64 encoding. For example:

FIGURE 115 User Data Text Block: Base64 Format

Name	Owner	Auto-Accept Policy
cloud-reg-01	JK	Accept-Policy-01


```

dGltZXpvcmluIklvIcm9wZS9Mb25kb24iCmFjY2VwdF9sawNlbnNlPSJZZXMiCmFj
Y2Vzc19rZX1fawQ9IkFLSUFJUVU0VEdPM1lCV0NQMK9RIgpszZWNYZXRfYWNjZXNz
X2t1eT0iMEozRlVMQi9nNmR4VlBoQ01QczBK525aQU4vYWFXy3prNETjwIthMSIK
cGFzc3dvcmluIjSZUNKTnlnYmciCm93bmVpPSJpd251ci0wV0FQLVo3VlotSUZ
NC1RUFhXIgpvd251cl9zZWNYZXQ9InBhc3N3b3JkIgpzZF9hZGRyZXNzPSIxMC42
Mi4xNjcuMjAxOjg5MDA5CnNkX2N1cnQ9Ik1JSUNXRENDQWVhZ0F3SUJBZ01KQ9t
S3UvS1FRcHduTUEwR0NTcUdTSWZzRFFkN3VUFNRVV4Q3pBSk5nT1ZCQV1UQWtG
Vk1STXdfUUV1EV1FRSURLcFRiMjFsTFZOMF1YUmxNU0V3SHdZRFZRUUtEQmhKYm5S
bGhtNWxkQ0JYYVdSbmFYUnpJRkIwZVNCWHRUXDIaGNOTVRVd05USTVNVFV6TORV
eVdoY05NVGN3TlRjNE1UVXpPRFV5V2pCRk1Rc3dDUV1EV1FRROV3SkJWVEVUTUJF
R0FCZ01UQ0YydlVTTSdCcm9wZS9Mb25kb24iCmFjY2VwdF9sawNlbnNlPSJZZXMiCmFj

```

☐ Show as text

 Copy to clipboard

- If either the **Owner** or **Auto-Accept Policy** fields are not specified in the summary entry for the Cloud Registration entry, you must enable the **Show as text** check box.

The lines relating to the unspecified **Owner** or the unspecified **Auto-Accept Policy** are then included with placeholder text that you can complete manually in the AWS system. See [“Creating the First vTM in a Cluster” on page 188](#).

- Click **Copy to Clipboard** to copy the displayed user data text block.

Once you have copied the user data text block, you can paste it directly into the AWS creation wizard, see [“Creating a Cloud-Based Virtual Traffic Manager” on page 188](#).

Registering an Externally-Deployed Virtual Traffic Manager

The Services Director VA enables you to manually register one or more externally-deployed vTM. This adds the vTM to the estate of the Services Director, from where it can be licensed, monitored and metered.

Note: You cannot manually register a vTM that uses vTM Communications Channel, including vTMs that are behind a NAT device. Instead, you must self-register the vTMs, see [“Self-Registering an Externally-Deployed Virtual Traffic Manager” on page 169](#).

You can register/license a vTM that is in a cluster. This process does not register other vTMs in the cluster, nor does it license them; you must independently register and license each node in a cluster.

Before you register an externally-deployed vTM, ensure that all required Services Director objects exist:

- The required Feature Pack. This lists the functions supported by the vTM, see [“Adding a Feature Pack to the Services Director” on page 129](#).
- The required Owner. This identifies the customer/owner for the vTM, see [“Adding an Owner to the Services Director” on page 141](#).
- The required Access Profile (optional). This identifies the authentication mechanism for the vTM, see [“Creating an Access Profile \(vTM User Authentication Only\)” on page 272](#).

The Services Director VA also enables you to deploy vTM. Each is deployed into an container using an existing instance host. The Services Director VA can then manage the lifecycle states of these vTMs, which is not supported for externally-deployed vTMs. For details, refer to the *Pulse Services Director Advanced User Guide*.

Preparing to Register a Virtual Traffic Manager (Universal FLA)

After you have completed the initial configuration of a Services Directors HA pair (see [“Preparing to Install the Services Director Virtual Appliance” on page 9](#), you can add one or more externally-deployed vTMs to the estate of the Services Director.

One method for achieving this is by manual registration of each vTM. Typically, these will use a Universal FLA License.

Note: You cannot manually register a vTM that uses vTM Communications Channel, including vTMs that are behind a NAT device. Instead, you must self-register the vTMs, see [“Self-Registering an Externally-Deployed Virtual Traffic Manager” on page 169](#).

You can register an externally-deployed vTM using a Universal FLA when:

- The vTM is installed and running.
- The vTM is at version 10.1 or later.
- You know the vTM’s hostname (in DNS-enabled networks) or IP address.
- The vTM’s REST API is enabled.

If any vTM is running an earlier version of the vTM software, or has its REST API disabled, you must manually install a Legacy FLA License onto the Services Director. See [“Preparing to Register a Virtual Traffic Manager \(Legacy FLA License\)” on page 160](#).

Note: To minimize delays in licensing, ensure that the clocks of your Services Directors and your vTMs are aligned.

Registering a Virtual Traffic Manager (Universal FLA)

The Services Director VA supports the registration and management of vTM instances from its **vTM Instances** page. After you have completed all initial setup operations, no vTM instances are registered.

Note: You can use this procedure to manually register an AWS vTM instance that has an elastic management IP address.

Note: You cannot manually register a vTM that is behind a NAT device. This process requires the vTM to be self-registered, see [“Self-Registering an Externally-Deployed Virtual Traffic Manager” on page 169](#).

If you wish to register a vTM whose REST API is disabled, see [“Registering a Virtual Traffic Manager \(Inactive REST API\)” on page 161](#).

Note: To minimize delays in licensing, ensure that the clocks of your Services Director(s) and your vTM instances are aligned.

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user.

The **Home** page appears.

3. Click the **Services** menu, and then click **Services Controller: vTM Instances**.

The **vTM Instances** page appears. After you have completed the Setup Wizard, this page contains no entries.

FIGURE 116 The vTM Instances Page: Before vTM Registrations

vTM Instances

► **Filters** Filtering by Lifecycle, Instance Health, Licensing Health

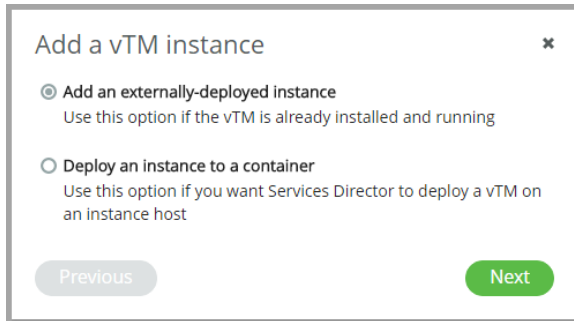
⊕ Add Show: 20 ▼ of 0 instances

Name ▾	License Name ▾	Bandwidth ▾	Feature Pack ▾	Version ▾	Cluster ▾	Instance Lifecycle ▾	Instance Health ▾	Licensing Health ▾
No Data								

- Click the plus symbol above the empty table.

If there is an instance host present on the Services Director, the following dialog box appears:

FIGURE 117 Adding an Instance Method



Add a vTM instance ✕

☒ **Add an externally-deployed instance**
Use this option if the vTM is already installed and running

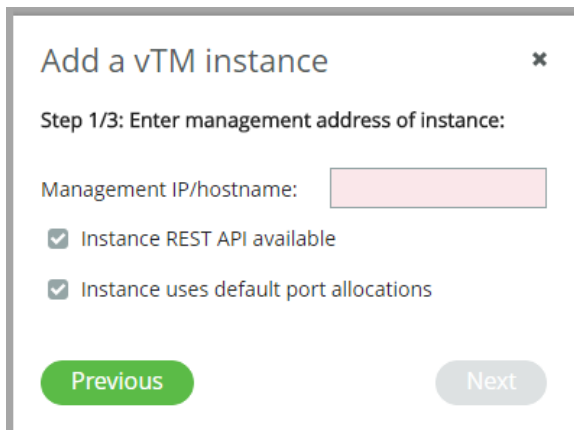
☐ **Deploy an instance to a container**
Use this option if you want Services Director to deploy a vTM on an instance host

Previous Next

- Click **Add an externally-deployed instance**, and then click **Next**.

After this (or if there is no instance host), a registration wizard appears:

FIGURE 118 Registration Wizard (1 of 3)



Add a vTM instance ✕

Step 1/3: Enter management address of instance:

Management IP/hostname:

☒ Instance REST API available

☒ Instance uses default port allocations

Previous Next

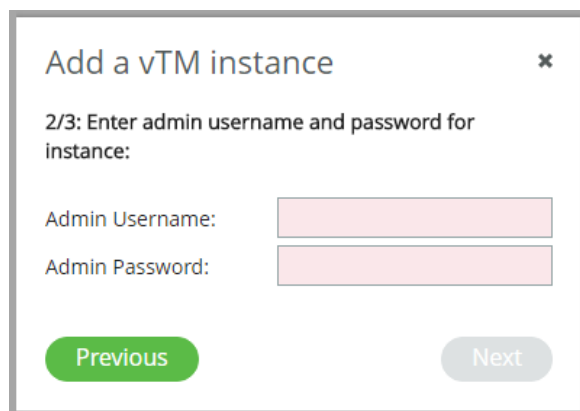
- Enter the hostname or IP address for the instance.

Note: From this wizard page, you can manually register an AWS vTM instance by specifying its elastic management IP address. In this instance, you must ensure that the AWS Security Groups for both the Services Director and the vTM are configured to support traffic flows, see [“Preparing an AWS Security Group” on page 40](#).

7. Click **Next**.

The next page of the wizard appears.

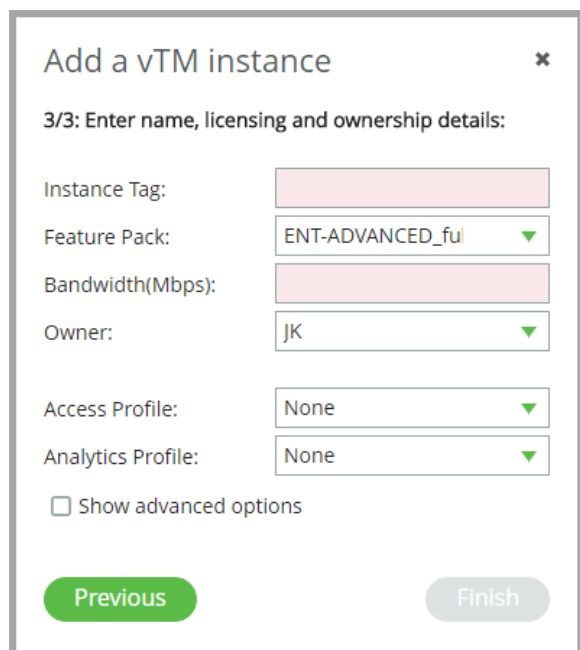
FIGURE 119 Registration Wizard (2 of 3)

A screenshot of the 'Add a vTM instance' wizard, step 2 of 3. The title bar says 'Add a vTM instance' with a close button. The main heading is '2/3: Enter admin username and password for instance:'. There are two input fields: 'Admin Username:' and 'Admin Password:'. At the bottom, there are two buttons: 'Previous' (green) and 'Next' (gray).

8. Enter the administration username and password, and click **Next**.

The next page of the wizard appears.

FIGURE 120 Registration Wizard (3 of 3)

A screenshot of the 'Add a vTM instance' wizard, step 3 of 3. The title bar says 'Add a vTM instance' with a close button. The main heading is '3/3: Enter name, licensing and ownership details:'. There are several input fields and dropdown menus: 'Instance Tag:', 'Feature Pack:' (with 'ENT-ADVANCED_full' selected), 'Bandwidth(Mbps):', 'Owner:' (with 'JK' selected), 'Access Profile:' (with 'None' selected), and 'Analytics Profile:' (with 'None' selected). There is a checkbox labeled 'Show advanced options'. At the bottom, there are two buttons: 'Previous' (green) and 'Finish' (gray).

9. Enter an **Instance Tag** for the vTM instance.

This is a user-facing name for the instance that will be used throughout the Services Director VA user interface. This tag can be changed at any time. It must be unique among non-deleted vTM instances registered on the Services Director, but can be reused as required.

That is, if an instance is deleted, its tag can be reused for a different instance.

10. Select a **Feature Pack** for the vTM instance.

This feature pack must be supported by your Services Director's License.

If the required Feature Pack is not defined on your Services Director, see [“Adding a Feature Pack to the Services Director” on page 129](#).

11. Enter a numeric **Bandwidth** (in Mbps) for the vTM instance.

This bandwidth must be available within your Services Director's Bandwidth License.

12. Either:

- Select an **Owner** for the vTM instance. See [“Adding an Owner to the Services Director” on page 141](#). OR
- Select *<create new>* from the **Owner** list, and type the name of a new owner. The Services Director will create a new Owner resource automatically when the registration completes. You can fully populate the Owner resource afterwards, see [“Viewing Full Details for an Owner” on page 143](#).

13. (Optional) Select an **Access Profile**.

This access profile identifies the authenticator and permission groups required for the user authentication on this vTM instance.

Note: Access profile is a cluster-level configuration property, and is typically set from the **vTM Cluster** page (see [“Creating a Virtual Traffic Manager Cluster” on page 233](#)). The current cluster-level setting is displayed in this dialogue. If you provide a new value for this property, the access profile will be applied to the vTM, *and all other vTM instances in its cluster*.

14. (Optional) Select an **Analytics Profile**.

This analytics profile identifies the vTM analytics settings for this vTM instance.

Note: Analytics profile is a cluster-level configuration property, and is typically set from the **vTM Cluster** page (see [“Creating a Virtual Traffic Manager Cluster” on page 233](#)). The current cluster-level setting is displayed in this dialogue. If you provide a new value for this property, the analytics profile will be applied to the vTM, *and all other vTM instances in its cluster*.

15. Click **Show advanced options** to view additional settings.

This access profile identifies the authenticator and permission groups required for the user authentication on this vTM instance.

Note: Access profile is a cluster-level configuration property, and is typically set on the vTM Cluster (see [“Creating a Virtual Traffic Manager Cluster” on page 233](#)). If selected, the access profile will be applied to the vTM, and all other vTM instances in its cluster.

FIGURE 121 Registration Wizard (3 of 3)

The **vTM Version** will automatically be the software version of your vTM.

16. Select the **License Name** of your Universal FLA License.

17. Click **Finish**.

The vTM is added to the **vTM Instances** table.

If this vTM is at version 10.1 or earlier, no cluster information is displayed.

If this vTM is at version 10.2 or later, its cluster is considered:

- If the vTM is in a cluster, the cluster is displayed as a Discovered cluster. The other vTMs in the cluster remain unregistered and unlicensed; you must independently register and license each node in a cluster.
- If this vTM is not in a cluster, a new cluster is created. This cluster has an automatically-generated name, and is a Discovered cluster.

See [“Working with Virtual Traffic Manager Clusters” on page 229](#).

FIGURE 122 First Added vTM Instance

vTM Instances

► Filters Filtering by Lifecycle, Instance Health, Licensing Health

+ Add Show: 20 ▼ of 1 instances

Name	License Name	Bandwidth	Feature Pack	Version	Cluster	Instance Lifecycle	Instance Health	Licensing Health
▶ cerulean-01	universal_v4	100	ENT-ADVANCED_full	19.1	Cluster-8D6X-VP0H-7S4A-FMYI	Active	OK	Licensed

This new entry shows basic details for the vTM instance. This includes a color-coded **Instance Lifecycle** status, **Instance Health** status, and **License Health** status. See [“Viewing Virtual Traffic Managers” on page 194](#).

The **Instance Health** status is supported on all vTMs at version 10.3 or later with a REST API enabled. Where it is not supported, it will be shown as *N/A*.

The **License Health** status will be *Pending* (blue) until the licensing is confirmed. This then changes to *Licensed* (green).

- Click the arrow to the left of the entry. The entry then expands to show the full details of the vTM instance.

FIGURE 123 Full Details for a vTM

▼ cerulean-01	universal_v4	100	ENT-ENTERPRISE_full	19.1	Cluster-CRCF-9WDA-T1HE-Z5WS	Active	OK	Licensed
---------------	--------------	-----	---------------------	------	-----------------------------	--------	----	----------

Instance Host Name: Host

Instance Name: cerulean-01

Bandwidth: 100 Mbps

Owner: JK ▼

License Name: universal_v4 ▼

Feature Pack: ENT-ENTERPRISE_ft ▼

vTM Management:

Admin Username: _servicesdirector

Admin Password: ••••••••

SNMP Address: 10.62. :161

REST Address: 10.62. :9070

REST API: Enabled ▼

UI Address: 10.62. :9090

vTM Servers:
Please click for more details (You will be redirected to the vTM's Diagnose page)

Apply Revert

Instance Type: Externally Deployed

Instance status:

Status: Active

Advanced Options:

vTM Cluster ID: Cluster-CRCF-9WDA-T1HE-Z5WS

UUID: 10758629-bf47-3701-

Certificate:

Extra Options: snmpcommunity=pu

On this detailed view:

- The **UUID** property is a unique identifier for the vTM. This property is only populated when the vTM registration request originates on the vTM.
- The **Certificate** property is only populated when the vTM Communications Channel feature is in use, see [“Working with vTM Communications Channel” on page 126](#).
- The **Extra Options** property lists advanced settings. For more information, refer to Configuration Options (config_options) in the *Pulse Services Director Advanced User Guide*.

19. Repeat this procedure to add other vTM instances.

FIGURE 124 Second Added vTM Instance

vTM Instances

► **Filters** Filtering by Lifecycle, Instance Health, Licensing Health

+ Add Show: 20 ▼ of 2 Instances

Name	License Name	Bandwidth	Feature Pack	Version	Cluster	Instance Lifecycle	Instance Health	Licensing Health
▶ cerulean-01	universal_v4	100	ENT-ADVANCED_full	19.1	Cluster-8D6X-VP0H-754A-FMYI	Active	OK	Licensed
▶ cerulean-02	universal_v4	100	ENT-ADVANCED_full	19.1	Cluster-8D6X-VP0H-754A-FMYI	Active	OK	Licensed

Preparing to Register a Virtual Traffic Manager (Legacy FLA License)

When you register an externally-deployed vTM, typically it is at version 10.1 (or later) and its REST API is enabled. See [“Registering a Virtual Traffic Manager \(Universal FLA\)” on page 154](#).

However, you can also add a vTM that has:

- A disabled REST API. See [“Registering a Virtual Traffic Manager \(Inactive REST API\)” on page 161](#).
- A software version of 10.0 (or earlier). See [“Registering a Virtual Traffic Manager \(Pre-10.1 vTM Software Version\)” on page 164](#).

You can register these vTM instances when:

- The vTM is installed and running.
- You know the management address for the vTM. The management address that you specify when registering the vTM should always match the hostname of the vTM being registered. That is:
 - If the vTM has been configured with a resolvable hostname, that same hostname should be used as the management address when registering.
 - If the vTM has been configured without a resolvable hostname (and an IP address used instead), that IP address should be used as the management address when registering.

Note: Where no DNS-system is configured, the use of hostnames should be avoided in the product.

- You have already installed a Legacy FLA License onto the Services Director. See [“Adding a Legacy FLA License to the Services Director” on page 144.](#)
- You have manually installed a Legacy FLA License onto the vTM. Refer to the manuals for the Pulse Secure Virtual Traffic Manager. This is not required when the REST API is active.

Pulse Secure recommends that you use vTM 10.1 or later and universal licensing wherever possible.

Registering a Virtual Traffic Manager (Inactive REST API)

The Services Director VA supports the registration and management of vTMs from its **vTM Instances** page. This process requires:

- A valid Legacy FLA License, keyed to the Service Endpoint Address of your Services Directors. If you do not have this, see [“Adding a Legacy FLA License to the Services Director” on page 144.](#)
- A Feature Pack that identifies the supported features for the vTM. If you do not have this, see [“Adding a Feature Pack to the Services Director” on page 129.](#)

Note: You cannot specify an access profile for a vTM when its REST API is disabled.

To register a vTM with an inactive REST API:

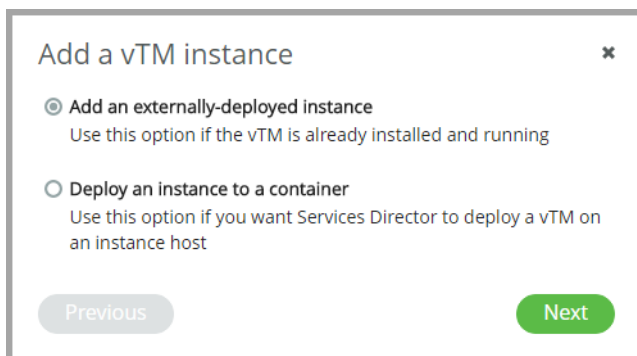
1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **Services Controller: vTM Instances**.

The **vTM Instances** page appears.

4. Click the plus symbol above the empty table.

If there is an instance host present on the Services Director, the following dialog box appears:

FIGURE 125 Adding an Instance Method



Add a vTM instance ✕

☒ **Add an externally-deployed instance**
Use this option if the vTM is already installed and running

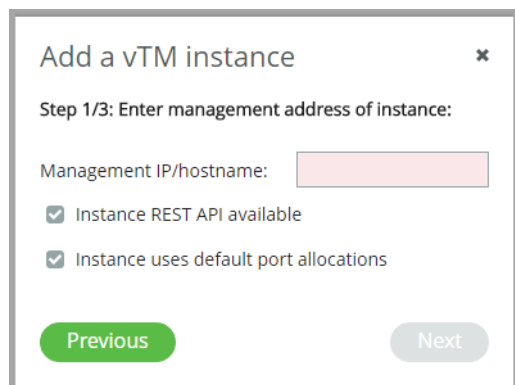
☐ **Deploy an instance to a container**
Use this option if you want Services Director to deploy a vTM on an instance host

Previous Next

Click **Add an externally-deployed instance**, and then click **Next**.

After this (or if there is no instance host), a registration wizard appears:

FIGURE 126 Registration Wizard (1 of 3)



The screenshot shows a dialog box titled "Add a vTM instance" with a close button (X) in the top right corner. Below the title, it says "Step 1/3: Enter management address of instance:". There is a text input field labeled "Management IP/hostname:". Below this field are two checked checkboxes: "Instance REST API available" and "Instance uses default port allocations". At the bottom, there are two buttons: "Previous" (green) and "Next" (grey).

5. Enter the management address for the vTM.

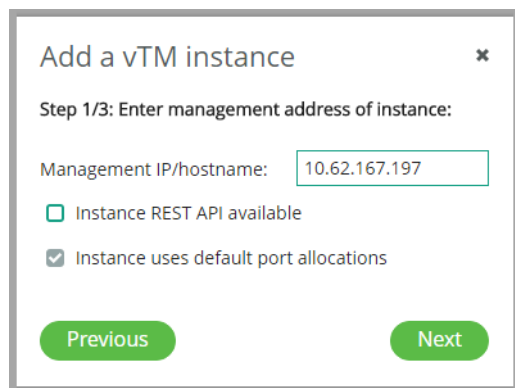
The management address that you specify when registering the vTM should always match the hostname of the vTM being registered. That is:

- If the vTM has been configured with a resolvable hostname, that same hostname should be used as the management address when registering.
- If the vTM has been configured without a resolvable hostname (and an IP address used instead), that IP address should be used as the management address when registering.

Where no DNS-system is configured, the use of hostnames should be avoided in the product.

6. Clear the **Instance REST API available** check box.

FIGURE 127 Clearing the Instance REST API Available Check Box



The screenshot shows the same dialog box as Figure 126, but with the "Instance REST API available" checkbox unchecked. The "Management IP/hostname:" field now contains the IP address "10.62.167.197". The "Previous" button is green, and the "Next" button is also green.

7. Click **Next**.

This option bypasses the second page of the wizard, and delivers you directly to the final page.

FIGURE 128 Registration Wizard (3 of 3)

8. Enter an **Instance Tag** for the vTM instance.

This is a user-facing name for the instance that will be used throughout the Services Director VA user interface. This tag can be changed at any time. It must be unique among non-deleted vTM instances registered on the Services Director, but can be reused as required.

That is, if an instance is deleted, its tag can be reused for a different instance.

9. Select a **Feature Pack** for the vTM instance.

This feature pack must be supported by your Services Director's License.

If the required Feature Pack is not defined on your Services Director, see [“Adding a Feature Pack to the Services Director” on page 129](#).

10. Enter a numeric **Bandwidth** (in Mbps) for the vTM instance.

This bandwidth must be available within your Services Director's Bandwidth License.

11. Select an **Owner** for the vTM instance. See [“Adding an Owner to the Services Director” on page 141](#).

Alternatively, select <create new> from the **Owner** list, and type the name of a new owner. The Services Director will create a new Owner resource automatically when the registration completes. You can fully populate the Owner resource afterwards, see [“Viewing Full Details for an Owner” on page 143](#).

12. Click **Finish**.

The vTM is added to the **vTM Instances** table.

The **Cluster** and software **Version** for this vTM are not shown, as the REST API is required to retrieve this information from the vTM.

If this vTM is not already in a cluster (and is at version 10.2 or later with the REST API enabled), a new cluster is created. This cluster has an automatically-generated name, and is a Discovered cluster. See [“Working with Virtual Traffic Manager Clusters” on page 229](#).

FIGURE 129 vTM Instance: Inactive REST API

vTM Instances

► Filters Filtering by Lifecycle, Instance Health, Licensing Health

⊕ Add Show: 20 ▼ of 3 instances

Name	License Name	Bandwidth	Feature Pack	Version	Cluster	Instance Lifecycle	Instance Health	Licensing Health
▶ cerulean-01	universal_v4	100	ENT-ADVANCED_full	17.3	Cluster-8D6X-VP0H-754A-FMYI	Active	OK	Licensed
▶ cerulean-02	universal_v4	100	ENT-ADVANCED_full	17.3	Cluster-8D6X-VP0H-754A-FMYI	Active	OK	Licensed
▶ viridian-01		50	ENT-ADVANCED_full			Active	N/A	Licensed

This new entry shows basic details for the vTM instance. This includes a color-coded **Instance Lifecycle** status, **Instance Health** status, and **License Health** status.

The **Instance Health** status is always *N/A* for vTMs using a Legacy FLA. This feature is only supported on vTMs at version 10.3 or later with a REST API enabled.

The **License Health** status will be *Pending* (blue) until the licensing is confirmed. This then changes to *Licensed* (green).

If the *Pending* status does not clear after a few minutes, log in to the affected vTM and investigate further.

Registering a Virtual Traffic Manager (Pre-10.1 vTM Software Version)

The Services Director VA supports the registration and management of vTM instances from its **vTM Instances** page. This process requires:

- A valid Legacy FLA License, keyed to the Service Endpoint Address of your Services Director instances. If you do not have this, see [“Adding a Legacy FLA License to the Services Director” on page 144](#).
 - A Feature Pack that identifies the supported features for the vTM. If you do not have this, see [“Adding a Feature Pack to the Services Director” on page 129](#).
1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
 2. Log in as the administration user. The **Home** page appears:

- Click the **Services** menu, and then click **Services Controller: vTM Instances**. The **vTM Instances** page appears.
- Click the plus symbol above the empty table.

If there is an instance host present on the Services Director, the following dialog box appears:

FIGURE 130 Adding an Instance Method

Add a vTM instance ✕

☒ **Add an externally-deployed instance**
Use this option if the vTM is already installed and running

☐ **Deploy an instance to a container**
Use this option if you want Services Director to deploy a vTM on an instance host

Previous Next

Click **Add an externally-deployed instance**, and then click **Next**.

After this (or if there is no instance host), a registration wizard appears:

FIGURE 131 Registration Wizard (1 of 3)

Add a vTM instance ✕

Step 1/3: Enter management address of instance:

Management IP/hostname:

☒ Instance REST API available

☒ Instance uses default port allocations

Previous Next

- Enter the management address for the vTM.

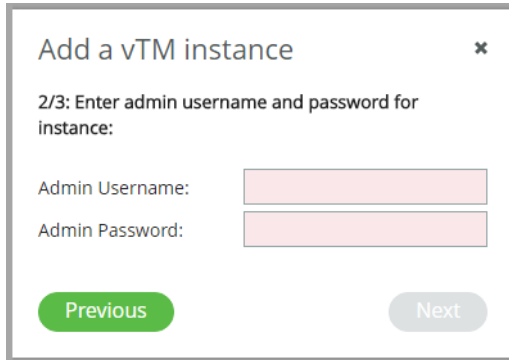
The management address that you specify when registering the vTM should always match the hostname of the vTM being registered. That is:

- If the vTM has been configured with a resolvable hostname, that same hostname should be used as the management address when registering.
- If the vTM has been configured without a resolvable hostname (and an IP address used instead), that IP address should be used as the management address when registering.

Where no DNS-system is configured, the use of hostnames should be avoided in the product.

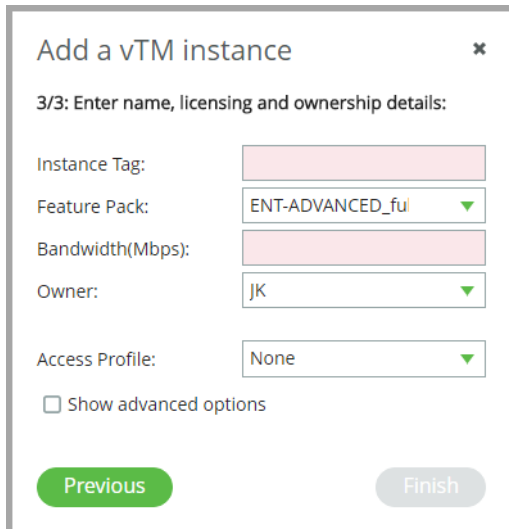
- Click **Next**. The next page of the wizard appears.

FIGURE 132 Registration Wizard (2 of 3)



- Enter the administration username and password.
- Click **Next**. The next page of the wizard appears.

FIGURE 133 Registration Wizard (3 of 3)



- Enter an **Instance Tag** for the vTM instance.

This is a user-facing name for the instance that will be used throughout the Services Director VA user interface. This tag can be changed at any time. It must be unique among non-deleted vTM instances registered on the Services Director, but can be reused as required.

That is, if an instance is deleted, its tag can be reused for a different instance.

- Select a **Feature Pack** for the vTM instance.

This feature pack must be supported by your Services Director's License.

If the required Feature Pack is not defined on your Services Director, see [“Adding a Feature Pack to the Services Director” on page 129](#).

11. Enter a numeric **Bandwidth** (in Mbps) for the vTM instance.

This bandwidth must be available within your Services Director's Bandwidth License.

12. Select an **Owner** for the vTM instance. See [“Adding an Owner to the Services Director” on page 141](#).

*Alternatively, select <create new> from the **Owner** list, and type the name of a new owner. The Services Director will create a new Owner resource automatically when the registration completes. You can fully populate the Owner resource afterwards, see [“Viewing Full Details for an Owner” on page 143](#).*

13. (Optional) Select an **Access Profile**.

This access profile identifies the authenticator and permission groups required for the user authentication on this vTM instance.

Note: Access profile is a cluster-level configuration property, and is typically set from the **vTM Cluster** page (see [“Creating a Virtual Traffic Manager Cluster” on page 233](#)). The current cluster-level setting is displayed in this dialogue. If you provide a new value for this property, the access profile will be applied to the vTM, *and all other vTM instances in its cluster*.

14. Click **Show advanced options** to view additional settings.

FIGURE 134 Registration Wizard (3 of 3)

Add a vTM instance ✕

3/3: Enter name, licensing and ownership details:

Instance Tag:

Feature Pack:

Bandwidth(Mbps):

Owner:

Access Profile:

☒ Show advanced options

vTM Version:

License Name:

The **vTM Version** will automatically be the software version of your vTM.

15. Select the **License Name** for your Legacy FLA License.

If the required Legacy FLA License is not listed, you must add it before you can register this vTM. See [“Adding a Legacy FLA License to the Services Director” on page 144](#).

16. Click **Finish**.

The vTM is added to the **vTM Instances** table.

The **Cluster** and software **Version** for this vTM are not shown, as version 10.2 and an active REST API are required to retrieve this information from the vTM.

FIGURE 135 vTM Instance: Pre-10.1 vTM Software Version

+ Add Show: 20 ▼ of 4 instances

Name ▾	License Name ▾	Bandwidth ▾	Feature Pack ▾	Version ▾	Cluster ▾	Instance Lifecycle ▾	Instance Health ▾	Licensing Health ▾
▶ cerulean-01	universal_v4	100	ENT-ADVANCED_full	17.3	Cluster-8D6X-VP0H-7S4A-FMYI	Active	OK	Licensed
▶ cerulean-02	universal_v4	100	ENT-ADVANCED_full	17.3	Cluster-8D6X-VP0H-7S4A-FMYI	Active	OK	Licensed
▶ viridian-01		50	ENT-ADVANCED_full			Active	N/A ▲	Licensed
▶ sunshine-01	legacy_9.3	50	ENT-ADVANCED_full	10.0		Active	N/A ▲	Licensed

This new entry shows basic details for the vTM instance. This includes a color-coded **Instance Lifecycle** status, **Instance Health** status and a **License Health** status. See [“Viewing Virtual Traffic Managers” on page 194](#).

The **Instance Health** status is always *N/A* for vTMs using a Legacy FLA. This feature is only supported on vTMs at version 10.3 or later with a REST API enabled.

The **License Health** status will be *Pending* (blue) until the licensing is confirmed. This then changes to *Licensed* (green).

If the *Pending* status does not clear after a few minutes, log in to the affected vTM and investigate further.

Self-Registering an Externally-Deployed Virtual Traffic Manager

The Services Director VA supports the self-registration of externally-deployed vTM. This adds vTMs to the estate of the Services Director, from where it can be licensed, monitored and metered.

This section describes the principles of vTM self-registration, and outlines the processing of self-registration requests on the Services Director.

Note: You must use self-registration for all vTMs that use the vTM Communications Channel, including vTMs that are behind a NAT device.

Overview: vTM Self-Registration (VMware)

After you have completed the initial configuration of the Services Director, you can add one or more externally-deployed vTMs to the estate of the Services Director.

One method for achieving this is by self-registration of the vTMs.

Note: Self-registration on the Services Director VA is also supported for cloud-based vTMs on AWS installations, see [“Overview: vTM Self-Registration \(Cloud\)” on page 186](#).

Note: Self-registration of vTMs that are in a private network behind a NAT requires the use of vTM Communications Channel on each vTM, see [“Working with vTM Communications Channel” on page 126](#).

Self-registration is initially configured from the vTM user interface. An Administrator configures the vTM so that it will request self-registration on a specified Services Director. Typically, this is done during the installation wizard for the vTM, see [“Requesting Self-Registration During vTM Installation” on page 172](#). However, this can also be done during later configuration of the vTM. See [“Requesting Self Registration on a Configured vTM” on page 177](#).

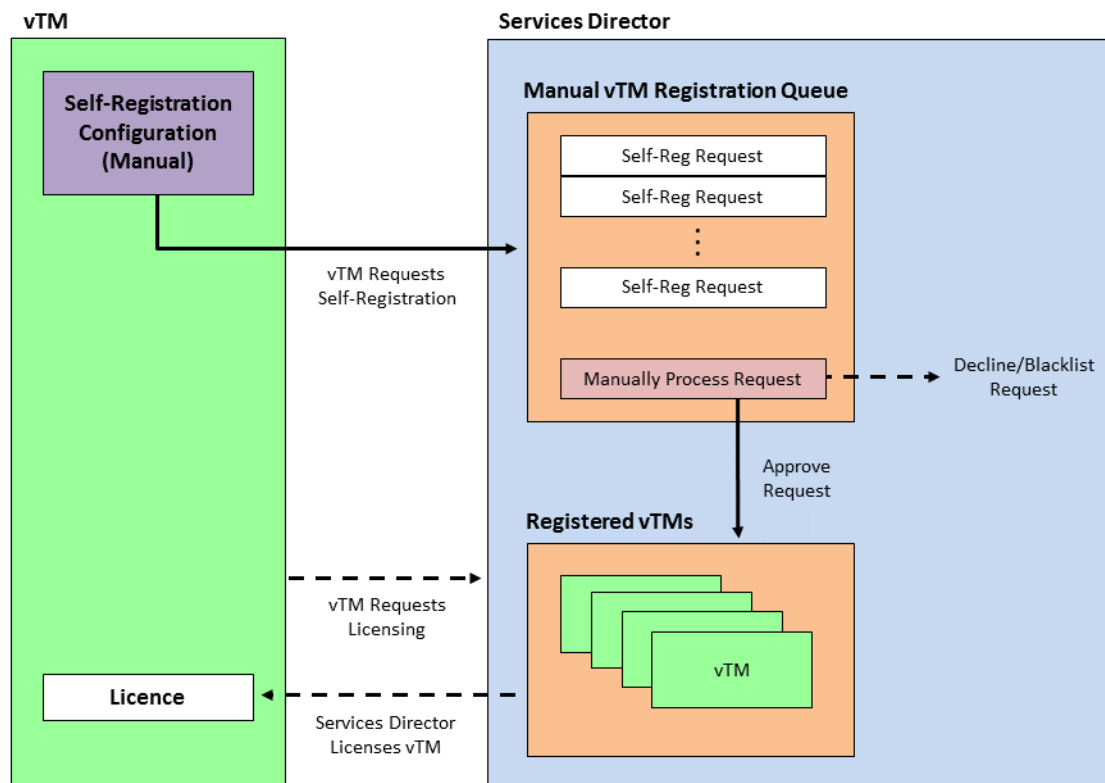
Self-registration can be either manual or automatic:

- Manual self-registration requires configuration of the vTM so that it requests self-registration on the Services Director.

When the request is received, the Services Director adds it to a queue of self-registration requests. The Administrator processes these manually as required, and can accept, decline or blacklist a request (see [“Processing Self-Registration Requests Manually” on page 182](#)).

Once a request is accepted, the vTM is added to the list of vTMs known to the Services Director. Licensing of the vTM can then occur as a separate process.

FIGURE 136 Manual Self-Registration of a vTM

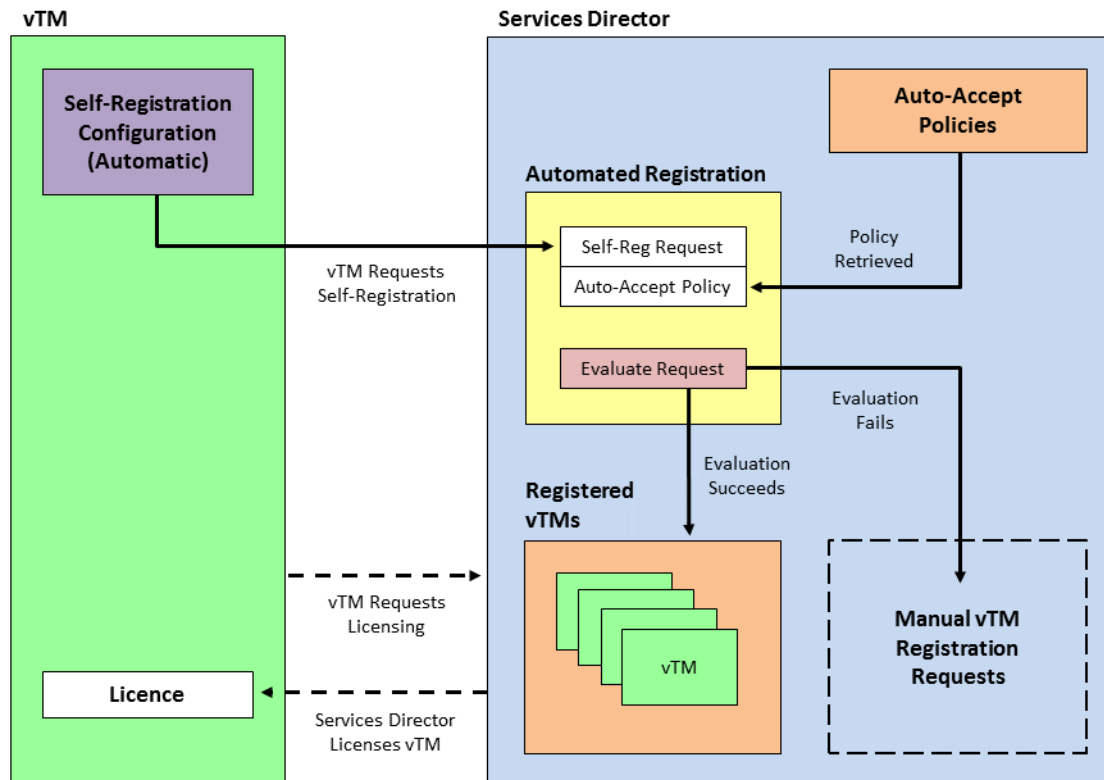


- Automatic self-registration requires configuration on both the vTM and the Services Director. An auto-accept policy must exist on the Services Director. This policy (one of many, potentially) defines the acceptance conditions and some fixed values for vTMs that use the policy. A policy must be referenced during the configuration of self-registration on the vTM.

When the request is received, the Services Director evaluates the request against the specified auto-accept policy, and will either accept or reject the vTM automatically.

Once accepted, the vTM is added to the list of vTMs known to the Services Director, and licensing of the vTM can then occur as a separate process. When rejected (for example, when there is insufficient bandwidth remaining, or the vTM is from outside the subnetwork), the vTM is added to the queue for manual self-registration requests instead, and the Administrator can process this in the usual way (see above).

FIGURE 137 Automatic Self-Registration of a vTM



Note: Once a vTM is configured for self-registration, it will make a self-registration request every time it restarts. The Services Director will assess this request, but will not process it if the vTM is already registered, blacklisted, or there is a pending self-registration request for the vTM.

Note: Once a self-registration request is received by the Services Director from a vTM, you must not change the cluster to which the vTM belongs until the registration request is accepted.

Requesting Self-Registration During vTM Installation

When you install the vTM VA, you can configure it for self-registration on the Services Director VA. Both manual and automatic self-registrations are supported.

Note: For a vTM at v19.1 or later, vTM Communications Channel (Comms Channel) is always enabled during the vTM's configuration wizard. To disable Comms Channel on an installed vTM at v19.1 or later, see [“Disabling Comms Channel on a vTM” on page 127](#).

Note: Once self-registration is requested by the vTM to the Services Director, you must not change the cluster to which a vTM belongs until the registration request is accepted.

Requesting Manual Self-Registration During the Installation of a vTM

This procedure enables you to configure a vTM for manual self-registration.

Note: For a vTM at v19.1 or later, vTM Communications Channel (Comms Channel) is always enabled during the vTM's configuration wizard. To disable Comms Channel on an installed vTM at v19.1 or later, see [“Disabling Comms Channel on a vTM” on page 127](#).

Note: For automatic self-registration, see [“Requesting Automatic Self-Registration During the Installation of a vTM” on page 174](#).

1. Install the vTM VA.
2. Log in to the vTM VA to start its installation wizard.
3. Progress through the Setup Wizard until the following page appears:

FIGURE 138 vTM Installation Wizard: License Key Page

Initial configuration, step 7 of 8

7. License Key

To use the traffic manager, you will need a valid license key. You have the following licensing options:

- ☒ Upload a license key for this traffic manager
- ☐ Register for flexible licensing using **Services Director**
- ☐ Skip licensing for now (traffic manager will run in **Developer mode** until licensing is configured)

Upload a new license key:

Key file: No file chosen

If you need to obtain a license key, please visit the **Brocade vTM website**.

4. Select **Register for flexible licensing using Services Director**. The page updates to include fields for self-registration:

FIGURE 139 vTM Installation Wizard: Requesting Self-Registration

Initial configuration, step 7 of 8

7. License Key

To use the traffic manager, you will need a valid license key. You have the following licensing options:

- ☐ Upload a license key for this traffic manager
- ☒ Register for flexible licensing using **Services Director**
- ☐ Skip licensing for now (traffic manager will run in **Developer mode** until licensing is configured)

This traffic manager will automatically register with your Services Director deployment.

Note: Services Director places some requirements on traffic managers it licenses in this way. If you proceed with this option:

- Services Director will be provided with user credentials for this traffic manager in order to install and configure licenses
- The REST API of this traffic manager will be enabled
- If this traffic manager is used as a template for other traffic manager appliances, these statements will be true for those as well

Services Director Address:

This should be the address of your Services Director's REST API, in the form <hostname/IP address>:<port>

Services Director Certificate:

You may provide details below to identify your registration request to the Services Director administrator.

Your e-mail address:

Registration Message:

Instance Owner:

Owner Secret:

Auto-accept Policy ID:

Advanced options

☐ This traffic manager appliance is for use as a template only (don't auto-register it with Services Director)

◀ Back

Next ▶

5. Specify the **Services Director Address**. This is the management address of the REST API port for the Services Director, as an <ip_address/host>:<port> pair.
6. Paste the Services Director's REST API SSL certificate as the **Services Director Certificate**. Contact the Services Director Administrator to obtain this.
7. (Optional) Specify **Your e-mail address**. If you provide this, the Services Director Administrator will receive a notification email when the self-registration request is received by the Services Director.

8. (Optional) Specify a **Registration Message**. This is seen by the Services Director Administrator when they view the self-registration request.
9. (Optional) Select an **Owner** for the vTM instance.

The owner entry was created in the Services Director, see [“Adding an Owner to the Services Director” on page 141](#).

10. Where you have selected an **Owner**, enter the **Owner Secret** password.

The password for the owner was created in the Services Director, see [“Adding an Owner to the Services Director” on page 141](#).

11. Do not enter an **Auto-accept Policy ID**. This is required for automatic self-registration only.
12. Ensure that the **Advanced Options** check box is clear. This is only required when creating a template vTM, see [“Working with vTM Templates” on page 277](#).
13. Click **Next** to go to the final wizard page and complete the wizard.

After the wizard completes, the vTM restarts.

The Services Director will receive a self-registration request from the vTM after the vTM restarts. The request is added to the queue of vTM self-registration requests, and can then be processed manually, see [“Accepting a Pending Self-Registration Request” on page 182](#).

Note: Once a vTM is configured for self-registration, it will make a self-registration request every time it restarts. The Services Director will assess this request, but will not process it if the vTM is already registered, or there is a *Pending* self-registration request for the vTM.

Note: Once a self-registration request is received by the Services Director from a vTM, you must not change the cluster to which the vTM belongs until the registration request is accepted.

Requesting Automatic Self-Registration During the Installation of a vTM

This procedure enables you to configure a vTM for automatic self-registration.

Note: For a vTM at v19.1 or later, vTM Communications Channel (Comms Channel) is always enabled during installation. To disable Comms Channel on an installed vTM at v19.1 or later, see [“Disabling Comms Channel on a vTM” on page 127](#).

Note: For manual self-registration, see [“Requesting Manual Self-Registration During the Installation of a vTM” on page 172](#).

1. Install the vTM VA.
2. Log in to the vTM VA to start its installation wizard.

3. Progress through the Setup Wizard until the following page appears:

FIGURE 140 vTM Installation Wizard: License Key Page

Initial configuration, step 7 of 8

7. License Key

To use the traffic manager, you will need a valid license key. You have the following licensing options:

- ☒ Upload a license key for this traffic manager
- ☐ Register for flexible licensing using **Services Director**
- ☐ Skip licensing for now (traffic manager will run in **Developer mode** until licensing is configured)

Upload a new license key:

Key file:

If you need to obtain a license key, please visit the **Brocade vTM website**.

4. Select **Register for flexible licensing using Services Director**. The page updates to include fields for self-registration:

FIGURE 141 vTM Installation Wizard: Requesting Self-Registration

Initial configuration, step 7 of 8

7. License Key

To use the traffic manager, you will need a valid license key. You have the following licensing options:

- ☐ Upload a license key for this traffic manager
- ☒ Register for flexible licensing using **Services Director**
- ☐ Skip licensing for now (traffic manager will run in **Developer mode** until licensing is configured)

This traffic manager will automatically register with your Services Director deployment.

Note: Services Director places some requirements on traffic managers it licenses in this way. If you proceed with this option:

- Services Director will be provided with user credentials for this traffic manager in order to install and configure licenses
- The REST API of this traffic manager will be enabled
- If this traffic manager is used as a template for other traffic manager appliances, these statements will be true for those as well

Services Director Address:

This should be the address of your Services Director's REST API, in the form <hostname/IP address> : <port>

Services Director Certificate:

You may provide details below to identify your registration request to the Services Director administrator.

Your e-mail address:

Registration Message:

Instance Owner:

Owner Secret:

Auto-accept Policy ID:

Advanced options

☐ This traffic manager appliance is for use as a template only (don't auto-register it with Services Director)

5. Specify the **Services Director Address**. This is the management address of the REST API port for the Services Director, as an <ip_address/host>:<port> pair.
6. Paste the Services Director's REST API SSL certificate as the **Services Director Certificate**. Contact the Services Director Administrator to obtain this.
7. (Optional) Specify **Your e-mail address**. If you provide this, the Services Director Administrator will receive a notification email when the self-registration request is received by the Services Director.
8. (Optional) Specify a **Registration Message**. This is seen by the Services Director Administrator when they view the self-registration request.
9. Select an **Owner** for the vTM instance. The owner entry was created in the Services Director, see ["Adding an Owner to the Services Director" on page 141](#).
10. Enter the **Owner Secret** password for the selected **Owner**. The password for the owner was created in the Services Director, see ["Adding an Owner to the Services Director" on page 141](#).
11. Enter the **Auto-accept Policy ID** of the auto-accept policy required for this vTM instance. The auto-accept policy was created in the Services Director, see ["Adding an Auto-Accept Policy to the Services Director" on page 147](#).
12. Ensure that the **Advanced Options** check box is clear. This is only required when creating a template vTM, see ["Working with vTM Templates" on page 277](#).
13. Click **Next** to go to the final wizard page and complete the wizard. After the wizard completes, the vTM restarts. The Services Director will receive a request for automatic self-registration the vTM after the vTM restarts. Either:
 - If the request can be processed automatically using the specified auto-accept policy, the vTM is added to the estate of the Services Director immediately, and subsequently licensed.
 - If the request cannot be processed automatically using the specified auto-accept policy, the request is added to the queue of vTM self-registration requests, and can then be processed manually, see ["Accepting a Pending Self-Registration Request" on page 182](#).

Once a vTM is configured for self-registration, it will make a self-registration request every time it restarts. The Services Director will assess this request, but will not process it if the vTM is already registered, or there is a *Pending* self-registration request for the vTM. note: Once a self-registration request is received by the Services Director from a vTM, you must not change the cluster to which the vTM belongs until the registration request is accepted.

Requesting Self Registration on a Configured vTM

You can configure an existing vTM to request self-registration.

To request self-registration on a configured vTM:

1. Log into the Services Director.
2. Click the **System** menu, and then click **Service SSL Certificate**.

The **Service SSL Certificate** page appears.

3. Click the **PEM** tab to view the SSL certificate in text form.
4. Copy the entire SSL certificate into your clipboard.
5. Log into the vTM.
6. Go to **System > Licenses**.

7. Under **Services Director Registration**:

- Set **remote_licensing!registration_server** to the public Services Director SEA and port. For example: *10.11.12.13: 8100*.
- Paste the Services Director SSL certificate from Step 4 into **remote_licensing!server_certificate**.
- Set **remote_licensing!owner** to the required Services Director Owner tag.
- Set **remote_licensing!owner_secret** to the secret/password for the Owner.
- (Optional) Set **remote_licensing!policy_id** to the UUID of the Services Director Self-Registration Policy. This is required for automatic self-registration only.
- Ensure that **remote_licensing!comm_channel_port** is set to *8102*.
- Set **remote_licensing!comm_channel_enabled** to the required value:
 - If it is set to *Yes*, Comms Channel will be enabled on the vTM.
 - If it is set to *No*, Comms Channel will be disabled on the vTM.

Note: The Comms Channel configuration on a vTM is not replicated to all vTMs in a cluster.

- (Optional) Set **remote_licensing!email_address** to an email address for system messages regarding the registration request.
 - (Optional) Set **remote_licensing!message** to a registration message that will be visible on the Services Director **vTM Instance Registrations** page.
8. Click **Save and Register**.

The vTM will register with the Services Director using the requested Comms Channel setting.

Note: For a vTM at v19.1 or later, vTM Communications Channel (Comms Channel) is always enabled during the vTM's configuration wizard. To disable Comms Channel on an installed vTM at v19.1 or later, see [“Disabling Comms Channel on a vTM” on page 127](#).

Viewing vTM Instance Registration Requests

The **vTM Instance Registrations** page lists all self-registration requests (both manual and automatic) that have been received by the Services Director from vTMs.

FIGURE 142 The vTM Instance Registration Page

vTM Instance Registrations

► Filters Showing Pending

Instance ID Info	Status	Registration Time	Email Address	Registration Message	Owner Validated?	Actions
▶ 10.62.169.171:9070	Pending	2016-10-03 15:35:22	admin@demo.com	Please register ...	✓	Accept Blacklist Decline
▶ 10.62.169.172:9070	Pending	2016-10-03 15:38:24			✗	Accept Blacklist Decline
▶ 10.62.169.173:9070	Pending	2016-10-03 15:38:54				Accept Blacklist Decline

◀ Page 1/1 ▶

See [“Understanding vTM Registration Requests” on page 179](#) for details of the headings.

You can **Accept**, **Blacklist** and **Decline** individual registrations from this list, see [“Processing Self-Registration Requests Manually” on page 182](#).

Expand a registration request to view its full details. For example:

FIGURE 143 vTM Instance Registration: Detailed View

Instance ID Info	Status	Registration Time	Email Address	Registration Message	Owner Validated?	Actions
▼ 10.62.169.171:9070	Pending	2016-10-03 15:35:22	admin@demo.com	Please register ...	✓	Accept Blacklist Decline
Registration ID : Reg-7LDJ-FZHF-XOVN-DDY9 Instance REST Address : 10.62.169.171:9070 Status : Pending Registration Time : 2016-10-03 15:35:22 Email Address : admin@demo.com Instance Version : 11.1a1 Owner : JK Registration Message : Please register automatically!						

This page also includes:

- A collapsed list of filters. These filters control which request state categories are displayed. See [“Filtering Self-Registration Requests” on page 181](#). Typically, you will view *Pending* requests only.

To view all requests for automatic self-registration, ensure you set the filter to include *Accepted* registrations.

- Paging controls for when there are larger numbers of registration requests.

Understanding vTM Registration Requests

Each entry in the table of vTM registration requests shows properties for a single self-registration request. Both automatic and manual self-registration requests are included. To view successful automatic self-registration requests, ensure that you have *Accepted* requests included, see [“Filtering Self-Registration Requests” on page 181](#).

Property	Description
Instance ID Info	<p>The information presented here depends on the use of vTM Communications Channel (Comms Channel):</p> <ul style="list-style-type: none"> • Where a registration request has come from a vTM that is using Comms Channel, the UUID of the vTM is displayed. • Where a registration request has come from a vTM that is not using Comms Channel, REST API address/port is displayed. <p>See “Working with vTM Communications Channel” on page 126.</p>
Status	The current state of the self-registration request. This determines the Actions that are supported for the request. See “Understanding Registration Status” on page 180 .
Registration Time	The time at which the Services Director received the self-registration request.
Email Address	The e-mail address of the administrator who configured the self-registration request on the vTM.
Registration Message	A text field. Typically, this will provide information for the Administrator who will process the self-registration request.
Owner Validated?	<p>Indicates whether owner information was received from the vTM, and whether it was valid:</p> <ul style="list-style-type: none"> • A tick indicates that owner/password information was received from the vTM, and that these have been validated against the Services Director's known owners. • A cross indicates that owner/password information was received from the vTM, but that it failed validation. • A blank column indicates that no owner/password information was received from the vTM.
Actions	A list of state transition actions that are valid from the current state. See “Understanding Registration Status” on page 180 .


Understanding Registration Status

The status of each self-registration request is displayed in the **vTM Instance Registration** page. See [“Viewing vTM Instance Registration Requests” on page 178](#).

Once self-registration is requested by the vTM to the Services Director, you must not change the cluster to which a vTM belongs until the registration request is accepted.

The lifecycle of a self-registration request is as follows:

FIGURE 144 State Model: Self-Registration Requests



When a self-registration request is received, it is given a *Pending* status.

For an automatic self-registration request, the auto-accept policy is then evaluated. Either:

- The evaluation of the auto-accept policy is successful. The request transitions automatically to *Accepted*, and the vTM is registered.
- The evaluation of the auto-accept policy is unsuccessful. The request retains its *Pending* status, and must then be resolved manually (see below).

For manual self-registration requests, you can transition it to:

- *Accepted*. You can manually transition a *Pending* request to *Accepted*, which completes the registration. See [“Accepting a Pending Self-Registration Request” on page 182](#).
- *Declined*. You can manually transition a *Pending* request to *Declined* if you do not wish to accept the request. See [“Declining a Pending Self-Registration Request” on page 183](#). You can transition a *Declined* request back to *Pending* if required.
- *Blacklisted*. You can manually transition a *Pending* request to *Blacklisted* if you do not wish to accept the request. See [“Blacklisting a Pending Self-Registration Request” on page 184](#). You can transition a *Blacklisted* request back to *Pending* if required.

A *Pending* request will transition to *Blacklisted* automatically after a defined timeout period. This defaults to 24 hours. See [“Updating Instance Registration Settings” on page 117](#).

The displayed states are subject to a status filter. By default, only *Pending* requests are shown. See [“Filtering Self-Registration Requests” on page 181](#).

To view automatic self-registration requests, you will need the *Accepted* requests to be visible.

Filtering Self-Registration Requests

You can filter the self-registration requests that are included on the **vTM Instance Registration** page. By default, only *Pending* requests are shown. When the filters are collapsed, a summary of the filter settings is shown:

FIGURE 145 vTM Self Registration Filters: Collapsed

► **Filters** Filtering by Pending, Blacklisted, Declined

Click the arrow on the left side of the filters to expand the **Status Filter** list.

FIGURE 146 vTM Self Registration Filters: Expanded

▼ **Filters** Filtering by Pending, Blacklisted, Declined

Status Filter

Pending ☒

Accepted ☐

Blacklisted ☒

Declined ☒

To view automatic self-registration requests that have been processed, the *Accepted* requests must be visible.

1. Click the **Catalogs** menu, and then click **Licensing: Instance Registrations**.

The **vTM Instance Registration** page appears.

2. Click the left arrow next to **Filters** to expand the **Status Filter** list.
3. Under **Status Filter**, select the check box for each required self-registration state.

Any state that is ticked is included in the table of self-registration requests.

Processing Self-Registration Requests Manually

All manual self-registrations and all failed automatic self-registrations are initially given a status of *Pending*. Each *Pending* request must be processed manually:

- [“Accepting a Pending Self-Registration Request” on page 182.](#)
- [“Declining a Pending Self-Registration Request” on page 183.](#)
- [“Blacklisting a Pending Self-Registration Request” on page 184.](#)
- [“Returning a Declined/Blacklisted Self-Registration Request to Pending” on page 185.](#)

Accepting a Pending Self-Registration Request

You can manually transition a *Pending* self-registration request to *Accepted*. You have the opportunity to review, change and confirm registration details before completing the process.

Once a vTM is registered, you cannot change the Accepted state of self-registration request.

1. Access your *Active Services Director* VA user interface from a browser, using the Service Endpoint Address of your Services Director.
2. Log in as the admin user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **Licensing: Instance Registrations**.

The **vTM Instance Registration** page appears.

4. Expand the filters, and ensure that *Pending* requests are included.
5. Locate the required *Pending* request.
6. Examine the information presented for the request, see [“Understanding vTM Registration Requests” on page 179.](#)

If additional information is required, expand the entry to view all details for the request, see [“Viewing vTM Instance Registration Requests” on page 178.](#)

7. In the **Actions** column for the request, click **Accept**.

The **Accept Registration** dialog box appears.

8. Enter an **Instance Name** for the vTM.

This is a user-facing name for the vTM that will be used throughout the Services Director VA user interface. This tag can be changed at any time. It must be unique among non-deleted vTM instances registered on the Services Director, but can be reused as required.

That is, if an instance is deleted, its tag can be reused for a different instance.

9. Enter an **Owner** for the vTM.

10. Select a **Feature Pack** for the vTM.

This feature pack must be supported by your Services Director's License. If the required Feature Pack is not defined on your Services Director, see [“Adding a Feature Pack to the Services Director” on page 129](#).

11. Enter a numeric **Bandwidth** (in Mbps) for the vTM.

This bandwidth must be available within your Services Director's Bandwidth License.

12. (Optional) Select an **Access Profile**.

This access profile identifies the authenticator and permission groups required for the user authentication on this vTM. See [“Working with User Authentication” on page 259](#).

13. Click **Accept**.

The state of the request changes to *Accepted*. The authenticator and permission groups in the access profile are applied to the vTM. Existing authenticators and permission groups may be overwritten, but none will be deleted. All members of a cluster are affected.

The vTM then appears as a registered vTM on the **vTM Instances** page.

Note: If the vTM uses Comms Channel, hyperlinks to the vTM will not be used, see [“Working with vTM Communications Channel” on page 126](#).

Declining a Pending Self-Registration Request

You can manually transition a *Pending* self-registration request to *Declined*. You can provide a reason for this decision if required.

You can exclude *Declined* requests from the **vTM Instance Registration** page if required by changing the Status Filter. See [“Filtering Self-Registration Requests” on page 181](#).

You can transition a *Declined* self-registration request back to *Pending*. See [“Returning a Declined/Blacklisted Self-Registration Request to Pending” on page 185](#).

1. *Active* Services Director VA user interface from a browser, using the Service Endpoint Address of your Services Director.
2. Log in as the admin user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **Licensing: Instance Registrations**.

The **vTM Instance Registration** page appears.

4. Expand the filters, and ensure that *Pending* requests are included.
5. Locate the required *Pending* request.

6. Examine the information presented for the request, see [“Understanding vTM Registration Requests” on page 179](#).

If additional information is required, expand the entry to view all details for the request, see [“Viewing vTM Instance Registration Requests” on page 178](#).

7. In the **Actions** column for the request, click **Decline**.

The **Decline Registration** dialog box appears.

8. (Optional) Enter your reasons for declining the request.

This information will be accessible to the vTM's Administrator.

9. Click **Decline** to close the dialog box. The state of the request changes to *Declined*.

Blacklisting a Pending Self-Registration Request

You can manually transition a *Pending* self-registration request to *Blacklisted*.

You can exclude *Blacklisted* requests from the **vTM Instance Registration** page if required by changing the Status Filter, see [“Filtering Self-Registration Requests” on page 181](#).

A *Pending* request will transition to *Blacklisted* automatically after a defined timeout period. This defaults to 24 hours. See [“Updating Instance Registration Settings” on page 117](#).

Note: You can transition a *Blacklisted* self-registration request back to *Pending*. See [“Returning a Declined/Blacklisted Self-Registration Request to Pending” on page 185](#).

1. Access the Services Director VA user interface from a browser, using the Service Endpoint Address of your Services Director.
2. Log in as the admin user.

The **Home** page appears.

3. Click the **Catalogs** menu, and then click **Licensing: Instance Registrations**.

The **vTM Instance Registration** page appears.

4. Expand the filters, and ensure that *Pending* requests are included.
5. Locate the required *Pending* request.
6. Examine the information presented for the request, see [“Understanding vTM Registration Requests” on page 179](#). If additional information is required, expand the entry to view all details for the request, see [“Viewing vTM Instance Registration Requests” on page 178](#).
7. In the **Actions** column for the request, click **Blacklist**.

The state of the request changes to *Blacklisted*.

Returning a Declined/Blacklisted Self-Registration Request to Pending

You can transition a *Declined/Blacklisted* self-registration request back to *Pending*. For example, you can choose to do this after an issue with a *Declined* request is resolved, or when a request that was *Blacklisted* automatically (see [“Updating Instance Registration Settings” on page 117](#)) still needs to be processed.

1. Active the Services Director VA user interface from a browser, using the Service Endpoint Address of your Services Director.
2. Log in as the admin user.

The **Home** page appears.

3. Click the **Catalogs** menu, and then click **Licensing: Instance Registrations**.

The **vTM Instance Registration** page appears.

4. Expand the filters, and ensure that *Declined/Blacklisted* requests are included.
5. Locate the required request.
6. In the **Actions** column for the request, click **Set to Pending**.

The state of the request changes to *Pending*.

Requesting Re-Registration of a vTM

After you have successfully self-registered a vTM, you may need to re-register it. For example, if the authorization credentials on the vTM change.

This process is performed entirely in the vTM user interface, under **System > Licenses > Services Director Registration**.

To force re-registration, update the registration details as required. Then, enable the **Force Re-Registration** check box and click **Save and Register**.

See the Virtual Traffic Manager documentation for full details of the vTM VA software.

Self-Registering a Cloud-Based Virtual Traffic Manager

The Services Director VA supports the automatic self-registration of cloud-based vTM instances. This adds cloud-based vTMs to the estate of the Services Director, from where it can be licensed, monitored and metered.

This section describes the principles of automatic self-registration for cloud-based vTMs.

Note: Self-registration of vTMs that are in a private network behind a NAT requires the use of vTM Communications Channel on each vTM, see [“Working with vTM Communications Channel” on page 126](#).

Overview: vTM Self-Registration (Cloud)

After you have completed the initial configuration of the Services Director, you can add one or more externally-deployed vTM to the estate of the Services Director.

One method for achieving this is by automatic self-registration a cloud-based vTM.

Note: Currently, cloud-based vTMs are supported on the Amazon Web Services (AWS) EC2 platform.

Note: Self-registration of vTMs that are in a private network behind a NAT requires the use of vTM Communications Channel on each vTM, see [“Working with vTM Communications Channel” on page 126](#).

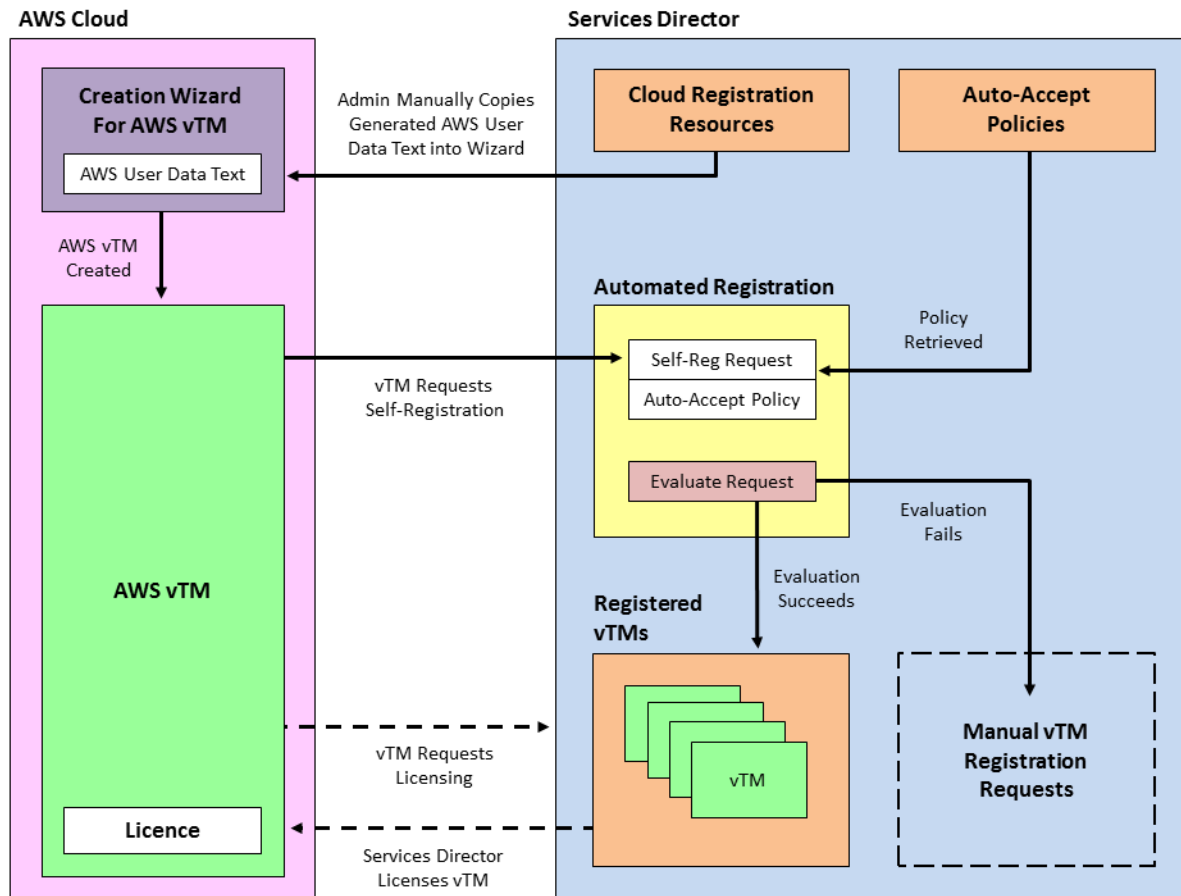
Cloud-based automatic registration begins on the Services Director, where a Cloud Registration resource must be created for one or more required deployments, see [“Adding a Cloud Registration Resource to the Services Director” on page 149](#). This resource identifies a number of properties that will be used by a cloud-based vTM, such as its Owner and the Self-Registration Policy that the Services Director will use to evaluate it.

Once a Cloud Registration resource has been created, a block of automatically-generated text becomes available on the Services Director. This text encapsulates the user data required by the AWS system to create the first cloud-based vTM in a cluster, and this vTM can automatically self-register on the Services Director. To do this, the administrator first manually copies this text into the AWS vTM creation wizard. Then, after the administrator specifies all other required network-specific details, the cloud-based AWS vTM is created. This process is described in [“Creating the First vTM in a Cluster” on page 188](#).

Self-registration of a cloud-based vTMs is intended to be automatic. The vTM makes a self-registration request to the Services Director. When the self-registration request is received, the Services Director evaluates the request against the specified self-registration policy, and will either accept or reject the vTM automatically.

When accepted, the vTM is added to the list of vTMs known to the Services Director. When rejected (for example, when there is insufficient bandwidth remaining, or the self-generated text does not include both an Owner and a Self-Registration Policy), the vTM is added to the queue of manual self-registration requests instead, and the Administrator can process manually, see [“Processing Self-Registration Requests Manually” on page 182](#).

FIGURE 147 Automatic Self-Registration of a Cloud-Based vTM



See **“Creating a Cloud-Based Virtual Traffic Manager” on page 188** for a full description of this process.

If you want to create additional cloud-based vTMs in the same cluster, you replace the user data text block for the Cloud Registration resource with the user data text block from the vTM's cluster, see **“Creating the Second vTM in a Cluster” on page 190**.

Once a self-registered vTM is known to the Services Director, the Services Director will respond to valid licensing requests by licensing the vTM, in the same way as for any other registered vTM.

Note: Once a vTM is configured for self-registration, it will make a self-registration request every time it restarts. The Services Director will assess this request, but will not process it if the vTM is already registered, blacklisted, or there is a pending self-registration request for the vTM.

Note: Once a self-registration request is received by the Services Director from a vTM, you must not change the cluster to which the vTM belongs until the registration request is accepted.

Note: A detailed description of the creation of an AWS cloud-based vTM can be found in the Virtual Traffic Manager documentation, refer to the *Pulse Virtual Traffic Manager Cloud Services Installation and Getting Started Guide*.

Creating a Cloud-Based Virtual Traffic Manager

You create one or more cloud-based vTM instances from the Amazon Web Services (AWS) system. To do this, you use a block of user data text that is created automatically by the Services Director, see **“Overview: vTM Self-Registration (Cloud)” on page 186** for details.

You must create each cloud-based instance individually. There are separate processes for:

- Creating the first cloud-based vTM in a cluster, see **“Creating the First vTM in a Cluster” on page 188**.
- Creating the second cloud-based vTM in a cluster, see **“Creating the Second vTM in a Cluster” on page 190**.
- All subsequent cloud-based vTMs in a cluster, see **“Creating Subsequent vTMs in a Cluster” on page 192**.

Creating the First vTM in a Cluster

The creation of a cloud-based vTM that is the first in a cluster is described below as a high-level process. Specific implementation choices will depend on your network configuration.

Before you perform this process, you must:

- Create the required Cloud Registration resource, see **“Adding a Cloud Registration Resource to the Services Director” on page 149**.
- Have the user data text block for this resource in your clipboard, see **“Viewing User Data Text for a Cloud Registration Resource” on page 151**.

Then, perform the following procedure.

1. On the Services Director, access the required Cloud Registration resource, and copy its user data text block to the clipboard. See **“Viewing User Data Text for a Cloud Registration Resource” on page 151**.
2. Access the Amazon Web Services (AWS) system and log in using your AWS credentials.
3. Access the EC2 dashboard.
4. Launch the process to create a new instance.

This starts a wizard that will lead you through the creation process.

5. On page 1 of the wizard (Choose AMI), locate and select the Amazon Machine Image (AMI) for the vTM from the AWS Marketplace.
6. On page 2 of the wizard (Choose Instance Type), select the required instance type.

7. On page 3 of the wizard (Configure Instance):

- Ensure the number of instances is 1. You can add more cloud-based instances to the cluster later, see [“Creating the Second vTM in a Cluster” on page 190](#).
- Select your network and subnetwork.
- You can choose to automatically assign a public IP for the new instance if required. By default, a public IP address is not assigned to a new instance. Your need to do this will depend on your specific networking configuration.
- Expand the advanced details, and paste in the AWS user data from your Cloud Registration resource.
- If your user data is plain text, add any incomplete properties, such as owner or auto-accept policy. If these are not specified, automatic self-registration will be unable to complete.

Note: If you do not intend to complete the owner or auto-accept policy properties, you must remove the incomplete entries from the pasted user data text block before continuing.

- Configure all other settings to your requirement.

8. On page 4 of the wizard (Add Storage), configure settings to match your network and requirement.

9. On page 5 of the wizard (Tag Instance), create a tag with **Key** set to “Name”, and **Value** set to the unique required name for your instance.

10. On page 6 of the wizard (Configure Security Group), either create a new security group, or select an existing one.

11. On page 7 of the wizard (Review):

- Review your choices and confirm. This effectively closes the wizard, but further configuration information is required.
- Create a new key pair. This key pair is used for this instance and all others that join its cluster.
- Download the key pair and save it in a safe location for future reference and use.
- Launch the instance.

The wizard closes and you are informed that the instance is being created.

Once the instance is created, it appears on the list of instances that is accessible from the EC2 dashboard.

When the Services Director receives the auto-registration request from the new cloud-based vTM, it will process the request:

- If automatic self-registration succeeds, the vTM will appear on the **vTM Instances** page, see [“Viewing Virtual Traffic Managers” on page 194](#). The vTM uses a new Discovered cluster. The name of the vTM is the private IP assigned by AWS to the vTM.

Note: If the vTM uses Comms Channel, hyperlinks to the vTM will not be used, see [“Working with vTM Communications Channel” on page 126](#).

- If automatic self-registration is unable to complete (for example, because of a missing owner or auto-accept policy), the registration request will appear as a *Pending* self-registration request on the **Instance Registrations** page. From there, you can manually process the request, see [“Processing Self-Registration Requests Manually” on page 182](#). Once you have accepted this self-registration request, you can create a second cloud-based vTM to the cluster, see [“Creating the Second vTM in a Cluster” on page 190](#).

Creating the Second vTM in a Cluster

The creation of a cloud-based vTM that is the second in a cluster is described below as a high-level process. Specific implementation choices will depend on your network configuration.

Note: Before you perform this process, you must create the first vTM in a cluster (refer to [“Creating the First vTM in a Cluster” on page 188](#)), and then access the user data text block from its vTM Cluster resource. This user data text block replaces the one that was used to create the first cloud-based vTM.

1. On the Services Director, access the vTM Cluster for the first vTM instance in the cluster, and copy its cluster text block to the clipboard. See [“Understanding Virtual Traffic Manager Cluster Details” on page 230](#).
2. Access the Amazon Web Services (AWS) system and log in using your AWS credentials.
3. Access the EC2 dashboard.
4. Launch the process to create a new instance.

This starts a wizard that will lead you through the creation process.

5. On page 1 of the wizard (Choose AMI), locate and select the Amazon Machine Image (AMI) for the vTM from the AWS Marketplace.
6. On page 2 of the wizard (Choose Instance Type), select the required instance type.
7. On page 3 of the wizard (Configure Instance):
 - Ensure the number of instances is 1. You can add more cloud-based instances to the cluster later, see [“Creating Subsequent vTMs in a Cluster” on page 192](#).
 - Select your network and subnetwork.

- You can choose to automatically assign a public IP for the new instance if required. By default, a public IP address is not assigned to a new instance. Your need to do this will depend on your specific networking configuration.
 - Expand the advanced details, and paste in the AWS user data from your vTM cluster.
 - Configure all other settings to your requirement.
8. On page 4 of the wizard (Add Storage), configure settings to match your network and requirement.
 9. On page 5 of the wizard (Tag Instance), enter a name for your instance.
 10. On page 6 of the wizard (Configure Security Group), select the existing security group that you used for the first instance in the cluster.
 11. On page 7 of the wizard (Review):
 - Review your choices and confirm. This effectively closes the wizard, but further configuration information is required.
 - Select the key pair that you created for the first vTM in the cluster. This key pair is used for all instances in the cluster.
 - Launch the instance.

The wizard closes and you are informed that the instance is being created.

Once the instance is created, it appears on the list of instances that is accessible from the EC2 dashboard.

When the Services Director receives the auto-registration request from the new cloud-based vTM, it will process the request:

- If successful, the vTM will appear on the **vTM Instances** page, see [“Viewing Virtual Traffic Managers” on page 194](#). This vTM shares its Discovered cluster with the first vTM in the cluster. The name of the vTM is the private IP assigned by AWS to the vTM.

Note: If the vTM uses Comms Channel, hyperlinks to the vTM will not be used, see [“Working with vTM Communications Channel” on page 126](#).
- If unsuccessful, the registration request will appear as a *Pending* self-registration request on the **Instance Registrations** page. From there, you can manually process the request, see [“Processing Self-Registration Requests Manually” on page 182](#). Once you have accepted this self-registration request, you can create additional cloud-based vTMs in the cluster, see [“Creating Subsequent vTMs in a Cluster” on page 192](#),

Creating Subsequent vTMs in a Cluster

Once you have created the first and second cloud-based vTMs in a cluster, creating additional vTMs in the cluster can be performed by duplicating the second vTM from the EC2 dashboard.

Note: You do not need to access and copy any user data text blocks during this process.

The creation of additional cloud-based vTMs in a cluster is described below as a high-level process. Specific implementation choices will depend on your network configuration.

1. Access the Amazon Web Services (AWS) system and log in using your AWS credentials.
2. Access the EC2 dashboard and view your instances.
3. Select the second instance in the cluster and issue a new action to create another instance like the one selected.

The instance creation wizard starts, and you are taken to page 7.

4. On page 7 of the wizard (Review):
 - Edit the tag for the new instance, so that it is unique. By default, it uses the same tag name as the duplicated instance.
 - Review your choices and confirm. This effectively closes the wizard, but further configuration information is required.
 - Select the key pair that you created for the first vTM in the cluster. This key pair is used for all instances in the cluster.
 - Launch the instance.

The wizard closes and you are informed that the instance is being created.

Once the instance is created, it appears on the list of instances that is accessible from the EC2 dashboard.

When the Services Director receives the auto-registration request from the new cloud-based vTM, it will process the request:

- If successful, the vTM will appear on the **vTM Instances** page, see [“Viewing Virtual Traffic Managers” on page 194](#). This vTM shares its Discovered cluster with the first vTM in the cluster. The name of the vTM is the private IP assigned by AWS to the vTM.
- If unsuccessful, the registration request will appear as a *Pending* self-registration request on the **Instance Registrations** page. From there, you can manually process the request, see [“Processing Self-Registration Requests Manually” on page 182](#).

Working with Virtual Traffic Managers

• Overview: Working with Virtual Traffic Managers.....	193
• Viewing Virtual Traffic Managers	194
• Viewing Full Details for a Virtual Traffic Manager	198
• Changing the Display Order of vTMs	199
• Filtering vTMs.....	200
• Updating Details for a Virtual Traffic Manager	201
• Understanding vServer Status	201
• Deleting a Virtual Traffic Manager	203
• Configuring Auto Cleanup of Virtual Traffic Managers.....	204
• Working with Application Templates (Enterprise Feature Tier)	207
• Relicensing Virtual Traffic Managers	221
• Processing Virtual Traffic Manager Metering Discrepancy Warnings	224

Overview: Working with Virtual Traffic Managers

Once you have installed your Pulse Secure Virtual Traffic Managers (vTMs), you manage them from the **vTM Instances** page of the Services Director VA. From this page, you can:

- View the basic status details for each vTM, including:
 - The lifecycle state of each vTM.
 - The instance health of each vTM.
 - The licensing health for each vTM.
- Show full details for each vTM.
- Change the order in which vTMs are displayed.
- Update the details for each vTM.
- Delete a vTM.
- Filter vTMs based on lifecycle state, instance health and licensing health.
- Change the lifecycle status for vTMs deployed from the Services Director.

Note: To register an externally-deployed vTM, see [“Adding Virtual Traffic Managers to the Services Director” on page 125](#).

Note: The operation of Traffic Management and Load Balancing on individual vTMs is not addressed by the Services Director product. This requires use of the Pulse Secure Virtual Traffic Manager software for each vTM.

Viewing Virtual Traffic Managers

The **vTM Instances** page shows a table of all vTM instances known by the Services Director.

This page also includes:

- A collapsed list of filters. These filters control which categories of vTM instances are displayed. See [“Filtering vTMs” on page 200](#).
- A count of instances.
- Paging controls for when there are larger numbers of vTM instances.

FIGURE 148 The vTM Instances Page

vTM Instances

► Filters Filtering by Lifecycle, Instance Health

+ Add

Show: 20 of 4 instances

Name	License Name	Bandwidth	Feature Pack	Version	Cluster	Instance Lifecycle	Instance Health	Licensing Health
▶ viridian-01	legacy_9.3	150	STM-400_full			Active	N/A	Licensed
▶ sunshine-01	legacy_9.3	200	STM-400_full	10.0		Active	N/A	Licensed
▶ violet-01	universal_v3	100	STM-400_full	10.3	Cluster-AC8L-ABCM-W5CR-ELSP	Active	OK	Licensed
▶ violet-02	universal_v3	100	STM-400_full	10.3	Cluster-RNPP-UIP9-RUA7-Q2JU	Active	OK	Licensed

◀ ◁ Page 1/1 ▷ ▶

Understanding Basic Details of a Virtual Traffic Manager

Each entry in the table of vTM instances shows basic details for the vTM.

Name	Description
Name	<p>The chosen name for the vTM. The name is displayed as a hyperlink, except where the vTM uses Comms Channel, see “Working with vTM Communications Channel” on page 126.</p> <p>Names can be edited, and reused after a vTM is deleted if required.</p>
License Name	The name of the FLA License for the vTM. This will either be a Universal FLA or a Legacy FLA, depending on the vTM settings.
Bandwidth	The maximum permitted bandwidth for this vTM (in Mbps).
Feature Pack	The chosen Feature Pack for the vTM.
Version	<p>The software version for the vTM.</p> <p>Where the vTM's REST API is unavailable, this is blank.</p>
Cluster	<p>The current cluster for the vTM. This is supported when:</p> <p>The vTM is deployed by the Services Director.</p> <p>The vTM is at version 10.2 or later with a REST API enabled.</p>

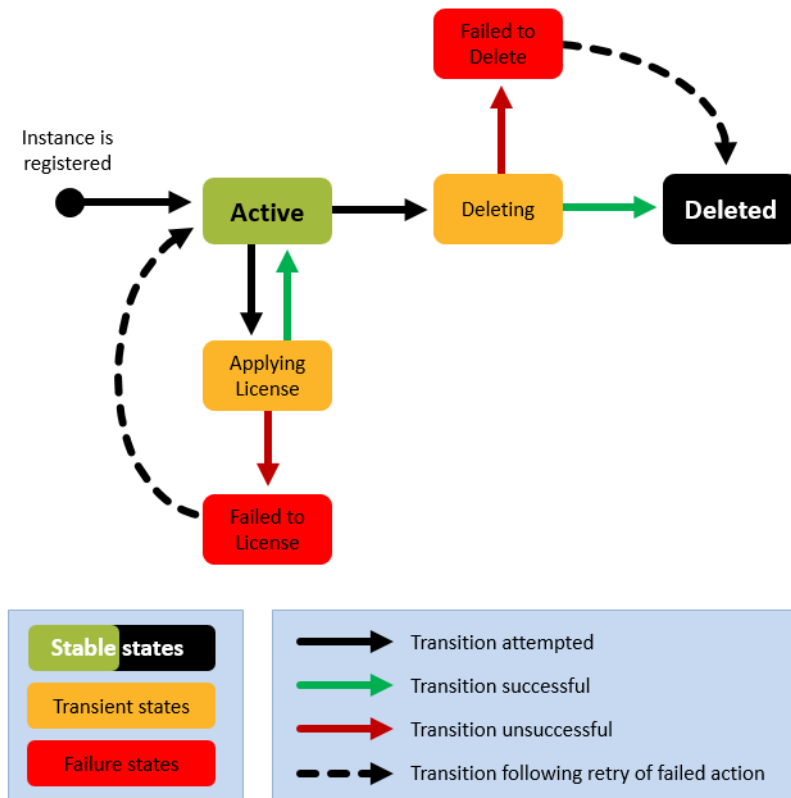
Name	Description
Instance Lifecycle	A colored indicator (green, blue, orange, red, black) and description of the vTM's lifecycle status. See “Understanding Lifecycle Status (Externally-Deployed vTMs)” on page 195 .
Instance Health	A colored indicator (green, blue, orange, red, black) and description of the vTM's current health status, which reflects the health of the cluster to which it belongs. See “Understanding the Instance Health of a Virtual Traffic Manager” on page 196 .
License Health	A colored indicator (green, blue, orange, red, black) and description of the vTM's current licensing health status. See “Understanding the Instance Health of a Virtual Traffic Manager” on page 196 .
Action	<p><i>Actions are only available for vTMs deployed by the Services Director.</i></p> <ul style="list-style-type: none"> When a vTM is Active, a Stop button is displayed. This enables you to stop the vTM, changing its status to <i>Idle</i>. A status of <i>Stopping</i> is displayed during this process. When a vTM is <i>Idle</i>, a Start button is displayed. This enables you to start the vTM, changing its status to <i>Active</i>. A status of <i>Starting</i> is displayed during this process.

Understanding Lifecycle Status (Externally-Deployed vTMs)

The **Instance Lifecycle** state of each vTM is displayed in the **vTM Instances** page.

When you register an externally-deployed vTM, the lifecycle operations supported by the Services Director VA are as follows:

FIGURE 149 Lifecycle States: Externally-Deployed vTMs



For most externally-deployed vTMs, the **Instance Lifecycle** state will remain *Active* until the vTM is deleted.

Note: The **Lifecycle Status** column for an externally-deployed vTM does *not* display a live monitoring status. As a result, if a vTM fails independently, this will not be indicated.

FIGURE 150 Lifecycle Status Column: Externally-Deployed vTMs

Active	A stable state
(Transitional)	A transitional state, indicating that an operation is in progress
Failed	A transitional state, indicating that an operation has failed
Deleted	A stable state

Note that:

- Where additional information is available, a red triangle is displayed next to the text of the status. Pause over this with your pointer to read the additional information in a dialog box.
- The displayed states are subject to the **Instance Status** filter. See [“Filtering vTMs” on page 200](#).

You can affect the **Lifecycle Status** of an externally-deployed vTM as follows:

- By deleting a vTM from its entry in the vTM table. See [“Deleting a Virtual Traffic Manager” on page 203](#).
- Other states are visible during relicensing.

Understanding Lifecycle Status (Deployed vTMs)

The **Instance Lifecycle** state of each vTM is displayed in the **vTM Instances** page.

When you deploy a vTM from the Services Director VA, it is deployed into a *container* on an instance host. This container enables full control of lifecycle operations for the vTM.

Refer to the *Pulse Services Director Advanced User Guide* for full details.

Understanding the Instance Health of a Virtual Traffic Manager

The **Instance Health** of each vTM is displayed in the **vTM Instances** page.

The displayed **Instance Health** of a vTM is a summary status that reflects the health of the *cluster* to which the vTM belongs. As a result, where cluster health is an issue, all vTMs in a cluster will typically display the same status.

Instance Health is reported as follows:

FIGURE 151 Instance Health Column

N/A	The cluster status on the instance is either not being monitored or the instance is deleted
N/A	The Traffic Manager uses an older software version (pre-v10.3) which does not support this feature
N/A	The cluster status on the instance is configured, but is not available. For example, when its REST API is not enabled
Okay	The cluster is in a <i>Healthy</i> state. This matches the dashboard state on the Traffic Manager itself.
Warning	The cluster is in a <i>Warning</i> state. This matches the dashboard state on the Traffic Manager itself.
Error	The cluster is in a <i>Serious</i> state. This matches the dashboard state on the Traffic Manager itself.

Note that:

- Instance health checks are only performed for vTMs at version 10.3 or later with an active REST API. For all other cases, the **Instance Health** is reported as *N/A*.
- Where additional information is available, a red triangle is displayed next to the text of the status. Pause over this with your pointer to read the additional information in a dialog box.
- The displayed states are subject to the **Instance Health** filter, see [“Filtering vTMs” on page 200](#).

Understanding the Licensing Health of a Virtual Traffic Manager

The **Licensing Health** of each vTM is displayed in the **vTM Instances** page.

The displayed **Licensing Health** of a vTM is a summary status, based on a number of licensing checks. Licensing is requested every three minutes using a callback mechanism. The method varies, depending on whether a Universal FLA or Legacy FLA License is in use on a vTM.

Licensing Health is reported as follows:

FIGURE 152 Licensing Health Column

Licensed	The Traffic Manager has called back and been licensed in the past three minutes
Pending	The Traffic Manager has been installed, but has not yet called back
Grace Period	The licensed Traffic Manager has not called back for over three minutes (but less than six weeks)
Expired	The licensed Traffic Manager has not called back for over six weeks (the grace period has expired)
N/A	Not Applicable: the Traffic Manager is not Active, and licensing checks are not required.

Note that:

- License checks are only performed for vTMs with an **Active Lifecycle Status**. For all other lifecycle states, the **Licensing Health** is reported as *N/A*.
- Where additional information is available, a red triangle is displayed next to the text of the status. Pause over this with your pointer to read the additional information in a dialog box.
- The displayed states are subject to the **Licensing Health** filter, see [“Filtering vTMs” on page 200](#).

Viewing Full Details for a Virtual Traffic Manager

The **vTM Instances** page shows a table of basic details for all vTM instances. To view full details for a vTM, click the arrow on the left side of the vTM's entry.

FIGURE 153 Viewing Full Details for a vTM Instance

Instance Name	Host	Bandwidth	CPU Usage	Owner	License Name	Feature Pack	Instance Type	Instance status	Licensing Health
scarlet-01	universal_v3	100	STM-400-Full	10.1	Active	N/A	Licensed		
violet-01	universal_v3	150	STM-400-Full	10.3b1	Violet-Cluster	Active	Error	Licensed	
scarlet-01	universal_v3	100	STM-400-Full	10.1	Active	N/A	Licensed		

violet-01 universal_v3 150 STM-400-Full 10.3b1 Violet-Cluster Active Error Licensed

Instance Host Name: Host Instance Type: Externally Deployed

Instance Name: violet-01 Instance status: Status: Active

Bandwidth: 150 Mbps Advanced Options: vTM Cluster ID: Violet-Cluster

CPU Usage: 0 Extra Options:

Owner: JK

License Name: universal_v3

Feature Pack: STM-400-Full

vTM Management:

Admin Username: admin

Admin Password: [REDACTED]

SNMP Address: 10.62.169.165.161

REST Address: 10.62.169.165.9070

REST API: Enabled

UI Address: 10.62.169.165.9090

vTM Servers:

Name	Pool	Port	Throughput(Mbps)	Vserver Health
VS-Pool-512	Pool-512	82	0	Error
VS-Pool-327	Pool-327	81	0	Warning
VS-Pool-421	Pool-421	80	0	OK

Please click for more details (You will be redirected to vTM Diagnose page)

Apply Revert

Note: The administration password for the vTM is not displayed by default. To reveal the administration password, click the eye button next to the **Password** field.

This view shows full details for the vTM, and includes a list of vServers with a status for each. See [“Understanding vServer Status” on page 201](#).

Changing the Display Order of vTMs

The **vTM Instances** page shows a table of all vTMs known by the Services Director.

The table of vTMs can be sorted according to any of the basic details, including **Lifecycle Status** and **Licensing Health** (see [“Understanding Basic Details of a Virtual Traffic Manager” on page 194](#)). For example, the table is sorted by default by ascending **Name**.

FIGURE 154 vTM Table Sorted By Ascending Name

Name ▾	License Name ▾	Bandwidth ▾	Feature Pack ▾	Version ▾	Cluster ▾	Instance Lifecycle ▾	Instance Health ▾	Licensing Health ▾
▶ sunshine-01	legacy_9.3	200	STM-400_full	10.0		Active	N/A	Licensed
▶ violet-01	universal_v3	100	STM-400_full	10.3	Cluster-AC8L-ABCM-W5CR-ELSP	Active	OK	Licensed
▶ violet-02	universal_v3	100	STM-400_full	10.3	Cluster-RNPP-UIP9-RUA7-Q2JU	Active	OK	Licensed
▶ viridian-01	legacy_9.3	150	STM-400_full			Active	N/A ▲	Licensed

To sort the table based on *ascending* values of any of the basic details, click the relevant column heading. For example, after clicking the **Bandwidth** heading, the same table is now sorted according to ascending **Bandwidth**.

FIGURE 155 vTM Table Sorted By Ascending Bandwidth

Name ▾	License Name ▾	Bandwidth ▴	Feature Pack ▾	Version ▾	Cluster ▾	Instance Lifecycle ▾	Instance Health ▾	Licensing Health ▾
▶ violet-02	universal_v3	100	STM-400_full	10.3	Cluster-RNPP-UIP9-RUA7-Q2JU	Active	OK	Licensed
▶ violet-01	universal_v3	100	STM-400_full	10.3	Cluster-AC8L-ABCM-W5CR-ELSP	Active	OK	Licensed
▶ viridian-01	legacy_9.3	150	STM-400_full			Active	N/A ▲	Licensed
▶ sunshine-01	legacy_9.3	200	STM-400_full	10.0		Active	N/A	Licensed

Clicking the column heading again will sort the table according to a *descending* view of the same basic detail. For example, after clicking the **Bandwidth** heading again, the same table is now sorted according to a descending value of **Bandwidth**.

FIGURE 156 vTM Table Sorted By Descending Bandwidth

Name ▾	License Name ▾	Bandwidth ▾	Feature Pack ▾	Version ▾	Cluster ▾	Instance Lifecycle ▾	Instance Health ▾	Licensing Health ▾
▶ sunshine-01	legacy_9.3	200	STM-400_full	10.0		Active	N/A	Licensed
▶ viridian-01	legacy_9.3	150	STM-400_full			Active	N/A ▲	Licensed
▶ violet-01	universal_v3	100	STM-400_full	10.3	Cluster-AC8L-ABCM-W5CR-ELSP	Active	OK	Licensed
▶ violet-02	universal_v3	100	STM-400_full	10.3	Cluster-RNPP-UIP9-RUA7-Q2JU	Active	OK	Licensed

Filtering vTMs

You can filter the vTM instances that are included on the **vTM Instances** page.

By default, the filters are collapsed, and a summary of filters is shown:

FIGURE 157 vTM Instance Filters: Collapsed

► **Filters** Filtering by Lifecycle, Instance Health

You can expand this to show the filters list.

FIGURE 158 vTM Instance Filters: Expanded

▼ **Filters** Filtering by Lifecycle, Instance Health

Basic Filters	Lifecycle Filter	Instance Health Filter	Licensing Health Filter	Cluster Filter
Name <input type="text"/>	Deleted <input type="checkbox"/>	N/A <input checked="" type="checkbox"/>	N/A <input checked="" type="checkbox"/>	N/A <input type="text" value="N/A"/>
	Active <input checked="" type="checkbox"/>	OK <input checked="" type="checkbox"/>	Licensed <input checked="" type="checkbox"/>	N/A
	Idle <input checked="" type="checkbox"/>	Warning <input checked="" type="checkbox"/>	Pending <input checked="" type="checkbox"/>	Cluster-RNPP-UIP9-RUA7-Q2JU
	Failed <input checked="" type="checkbox"/>	Error <input checked="" type="checkbox"/>	Warning <input checked="" type="checkbox"/>	Cluster-AC8L-ABCM-W5CR-ELSP
			Failed <input checked="" type="checkbox"/>	

The following filters are supported, which can be used in combination:

- **Basic Filters** - this filters vTMs by name. This supports *regular expressions* for search purposes.
- **Lifecycle Filter** - this filters vTMs by instance lifecycle status. Any of the four lifecycle states can be included/excluded. That is: *Active*, *Idle*, *Failed*, *Deleted*. You cannot filter using any of the (orange) supported transitional states.

vTMs with the *Deleted* instance lifecycle state are not included by default.
- **Instance Health Filter** - this filters vTMs by license health. Any of the four licensing states can be included/excluded. That is: *Error*, *Warning*, *OK* or *N/A*.
- **Licensing Health Filter** - this filters vTMs by license health. Any of the licensing states can be included/excluded. That is: *Licensed*, *Pending*, *Warning*, *Failed* or *N/A*.
- **Cluster Filter** - this filters vTMs using a single selected cluster. The list of clusters includes both Discovered and User Created clusters, see [“Working with Virtual Traffic Manager Clusters” on page 229](#).

Perform the following procedure:

1. Click the **Services** menu, and then click **Services Director: vTM Instances**.

The **vTM Instances** page appears.

2. Under **Basic Filters**, type a **Name** if required. This supports *regular expressions* for search purposes. This filter is applied automatically as you type.

When a **Name** filter is set the summary of filters includes "Name".

3. Under **Lifecycle Status**, select the check box for each required instance lifecycle state.

Any state that is ticked is included in the table of vTMs.

Deleted vTMs are not included by default. To include these, select the **Deleted** check box.

4. Under **Instance Health**, select the check box for any required instance health states.

Any state that is ticked is included in the table of vTMs.

5. Under **License Health**, select the check box for any required licensing health states.

Any state that is ticked is included in the table of vTMs.

6. Under **Cluster**, select the required cluster from the drop-down list.

The table of vTMs is limited to vTMs that are in the selected cluster.

Updating Details for a Virtual Traffic Manager

You can update many of the details of a vTM from the **vTM Instances** page.

1. Click the **Services** menu, and then click **Services Director: vTM Instances**. The **vTM Instances** page appears.
2. Locate the vTM's entry in the table of vTMs.
3. Click the arrow on the left side of the vTM's entry. The entry expands to show full details for the vTM.
4. Make the required changes to the vTM's details.
5. Click **Apply**.

Understanding vServer Status

Each vTM will have one or more vServers. Each vServer is responsible for balancing incoming traffic across a pool of nodes, as configured on the Pulse Secure Virtual Traffic Manager itself.

A list of vServers is included in the vTM detailed view on the **vTM Instances** page. The vTM must be at version 10.3 or later with the REST API available. For example:

FIGURE 159 Virtual Traffic Manager Details: vServers

violet-01
universal_v3
150
STM-400-Full
10.3
Violet-Cluster
Active
Error
Licensed

Instance Host Name: Host
Instance Name: violet-01
Bandwidth: 150 Mbps
CPU Usage: 0
Owner: JK
License Name: universal_v3
Feature Pack: STM-400-Full

Instance Type: Externally Deployed
Instance status:
Status: Active
Advanced Options:
vTM Cluster ID: Violet-Cluster
Extra Options:

vTM Management:
Admin Username: admin
Admin Password:
SNMP Address: 10.62.169.165.161
REST Address: 10.62.169.165.9070
REST API: Enabled
UI Address: 10.62.169.165.9090

vTM Servers:

Name	Pool	Port	Throughput(Mbps)	Vserver Health
VS-Pool-512	Pool-512	82	0	Error
VS-Pool-327	Pool-327	81	0	Warning
VS-Pool-421	Pool-421	80	0	OK

Please click for more details (You will be redirected to vTM Diagnose page)

Apply
Revert

In this example, the **vTM Servers** list shows three vServers:

- *VS-Pool-512* is in an *Error* state. This indicates that all of its nodes are in error. Pausing the pointer over the warning triangle will list failed pool nodes.
- *VS-Pool-327* is in a *Warning* state. This indicates that some (but not all) of its nodes are in error. Pausing the pointer over the warning triangle will list failed pool nodes.
- *VS-Pool-421* is in an *OK* state. This indicates that all of the vServer pool nodes are working.

The **vTM Servers** list is limited to ten vServers, but by default this list displays in descending order of severity. That is, vServers showing an *Error* at the top, then vServers showing warnings, then vServers with no errors.

Note: To investigate any listed errors, click the **Please click for more details** control. You will be redirected to **vTM Diagnose** page on the vTM software, outside of the Services Director VA.

Deleting a Virtual Traffic Manager

You can delete a vTM from the **vTM Instances** page.

When you delete an externally-deployed vTM:

- The vTM itself is not actually deleted. It continues to exist, and remains registered. However, monitoring, metering and licensing checks for the vTM are halted.
- The **Lifecycle Status** of the vTM changes to *Deleted*.
- The **Licensing Health** of the vTM changes to *N/A*.
- The **Name** of a *Deleted* vTM can be reused by a different vTM.

Note: vTMs with the *Deleted* state are not included in the default filter settings for the **vTM Instances** page. To include these vTMs in the **vTM Instances** page, see [“Filtering vTMs” on page 200](#).


When you delete a vTM that was deployed by the Services Director VA:

- The vTM must be in an *Idle* state.
- The vTM itself is deleted.
- The vTM’s container is deleted.
- The **Lifecycle Status** of the vTM changes to *Deleted*.
- The **Instance Health** of the vTM changes to *N/A*.
- The **Licensing Health** of the vTM changes to *N/A*.
- The **Name** of a *Deleted* vTM can be reused by a different vTM.

To delete a vTM:

1. Click the **Services** menu, and then click **Services Director: vTM Instances**. The **vTM Instances** page appears.
2. Locate the vTM’s entry in the table of vTMs.
3. To the right of the vTM’s entry, click the **X** control. A confirmation control appears.

FIGURE 160 The Delete Confirmation Control



4. Click **Delete**.

Configuring Auto Cleanup of Virtual Traffic Managers

You can configure Services Director to automatically delete registered vTM instances that have failed. This may be under specific circumstances, such as when a vTM is used to perform a transient service, and that service has ended.

Note: The deletion of a vTM from Services Director does not delete the vTM itself.

There are two configurations supported:

- Services Director deletes any *automatically self-registered* vTM that has failed. For details of this registration process, see [“Self-Registering an Externally-Deployed Virtual Traffic Manager” on page 169](#).
- Services Director deletes *any* vTM that has failed.

Both configurations require configuration of all **Instance Failure Period** and **Instance Monitor Interval** settings in the Services Director General Settings, see [“Updating Monitoring Settings” on page 115](#).

To configure automatic deletion of failed vTMs:

1. Access the **System > General Settings** page.
2. Under **Monitoring**, update the following settings:
 - **Instance Monitor Interval** - the length of the *monitoring cycle*. That is, the period of time, in seconds, between each Services Director attempt to retrieve monitoring information from each vTM. The default value is 60. When vTM monitoring information cannot be retrieved by Services Director for this period, the **Instance Health** of a vTM instance will change to *Error* on the **vTM Instances** page.
 - **Instance Failure Period** - the period of time, in seconds, after which the instance is considered to have failed if vTM monitoring information cannot be retrieved by Services Director. The default value is 180. When the vTM fails, auto-deletion will be triggered on eligible vTMs.

Note: Typically, the **Instance Failure Period** will be several times longer than the **Instance Monitor Interval**.

3. Under **Auto Cleanup vTMs**, choose the required setting:
 - To delete only automatically self-registered vTMs that fail, click **Self Registered Auto Accepted**.

The **Auto Cleanup vTMs Status** changes to *Self Registered Auto-Accepted*.

FIGURE 161 Configure Auto Cleanup of Failed Auto Self-Registered vTMs

Auto Cleanup vTMs

Auto Cleanup vTMs Status: **Self Registered Auto Accepted** Services Director can automatically mark instances as deleted if they repeatedly fail monitoring

Off Self Registered Auto Accepted All

- To delete all vTMs that fail, click **All**.

The **Auto Cleanup vTMs Status** changes to *All*.

FIGURE 162 Configure Auto Cleanup of Failed vTMs

Auto Cleanup vTMs

Auto Cleanup vTMs Status:

All vTMs

Services Director can automatically mark instances as deleted if they repeatedly fail monitoring

Off

Self Registered Auto Accepted

All

- (Optional) To disable Auto Cleanup, click **Off**.

The **Auto Cleanup vTMs Status** changes to *Off*.

Once the configuration process is complete, vTMs of the selected type will be deleted from the Services Director in the event of a vTM failure. See also [“Example of Auto Cleanup” on page 205](#).

Example of Auto Cleanup

In the following example:

- The vTMs *vermillion-01* and *vermillion-02* were manually registered on Services Director.
- The vTMs *cerulean-01* and *cerulean-02* were automatically self-registered on Services Director.
- Auto Cleanup is configured so that *automatically self-registered* vTMs will be automatically deleted from Services Director in the event of failure.

FIGURE 163 Example: vTM Instances Page

vTM Instances

Filters: Filtering by Lifecycle, Instance Health, Licensing Health

Add

Show: 20 of 4 instances

Name	License Name	Bandwidth	Feature Pack	Version	Cluster	Instance Lifecycle	Instance Health	Licensing Health
▶ vermillion-01	universal_v4	50	ENT-ENTERPRISE_full	19.1	Cluster-93AW-HPKU-H9B6-OJOV	Active	OK	Licensed
▶ vermillion-02	universal_v4	50	ENT-ENTERPRISE_full	19.1	Cluster-93AW-HPKU-H9B6-OJOV	Active	OK	Licensed
▶ cerulean-01	universal_v4	80	ENT-ENTERPRISE_full	19.1	Cluster-NBKN-NJZ5-HDOB-BG2N	Active	OK	Licensed
▶ cerulean-02	universal_v4	100	ENT-ENTERPRISE_full	19.1	Cluster-NBKN-NJZ5-HDOB-BG2N	Active	OK	Licensed

After a monitoring cycle, if monitoring information cannot be retrieved from *Vermillion-01* and *Cerulean-01*, their **Instance Health** and **Licensing Health** update to indicate this.

FIGURE 164 Example: Monitoring Fails on vTMs

vTM Instances

Filters Filtering by Lifecycle, Instance Health, Licensing Health

➕ Add Show: 20 of 4 instances

Name	License Name	Bandwidth	Feature Pack	Version	Cluster	Instance Lifecycle	Instance Health	Licensing Health
▶ vermillion-01	universal_v4	50	ENT-ENTERPRISE_full	19.1	Cluster-93AW-HPKU-H9B6-OJOV	Active	Error	Grace period
▶ vermillion-02	universal_v4	50	ENT-ENTERPRISE_full	19.1	Cluster-93AW-HPKU-H9B6-OJOV	Active	OK	Licensed
▶ cerulean-01	universal_v4	80	ENT-ENTERPRISE_full	19.1	Cluster-NBKN-NJZ5-HDOB-BG2N	Active	Error	Grace period
▶ cerulean-02	universal_v4	100	ENT-ENTERPRISE_full	19.1	Cluster-NBKN-NJZ5-HDOB-BG2N	Active	OK	Licensed

Once the failure period is reached without monitoring information being retrieved, auto cleanup triggers:

- *Vermilion-01* is not deleted, as it was *not* automatically self-registered.
- *Cerulean-01* is deleted, as it was automatically self-registered.

FIGURE 165 Example: Automatic Self-Registered vTM Deleted

vTM Instances

Filters Filtering by Lifecycle, Instance Health, Licensing Health

➕ Add Show: 20 of 3 instances

Name	License Name	Bandwidth	Feature Pack	Version	Cluster	Instance Lifecycle	Instance Health	Licensing Health
▶ vermillion-01	universal_v4	50	ENT-ENTERPRISE_full	19.1	Cluster-93AW-HPKU-H9B6-OJOV	Active	Error	Grace period
▶ vermillion-02	universal_v4	50	ENT-ENTERPRISE_full	19.1	Cluster-93AW-HPKU-H9B6-OJOV	Active	OK	Licensed
▶ cerulean-02	universal_v4	100	ENT-ENTERPRISE_full	19.1	Cluster-NBKN-NJZ5-HDOB-BG2N	Active	OK	Licensed

To confirm the deletion, expand the **Filters** and view *Deleted* vTMs to see the deleted *Cerulean-01* vTM:

FIGURE 166 Example: Viewing Deleted vTMs

vTM Instances

Filters Filtering by Lifecycle, Instance Health, Licensing Health

Basic Filters Name

Lifecycle Filter	Instance Health Filter	Licensing Health Filter	Cluster Filter
Deleted <input checked="" type="checkbox"/>	N/A <input checked="" type="checkbox"/>	N/A <input checked="" type="checkbox"/>	N/A
Active <input type="checkbox"/>	OK <input checked="" type="checkbox"/>	Licensed <input checked="" type="checkbox"/>	
Idle <input type="checkbox"/>	Warning <input checked="" type="checkbox"/>	Pending <input checked="" type="checkbox"/>	
Failed <input type="checkbox"/>	Error <input checked="" type="checkbox"/>	Warning <input checked="" type="checkbox"/>	
		Failed <input checked="" type="checkbox"/>	

➕ Add Show: 20 of 1 instances

Name	License Name	Bandwidth	Feature Pack	Version	Cluster	Instance Lifecycle	Instance Health	Licensing Health
▶ cerulean-01	universal_v4	80	ENT-ENTERPRISE_full		Cluster-NBKN-NJZ5-HDOB-BG2N	Deleted	N/A	N/A

Note: Deleted vTMs are purged from Services Director after a default period of 42 days. This period can only be set from the REST API, see the *Services Director Advanced User Guide*.

Working with Application Templates (Enterprise Feature Tier)

Once a clustered vTM is registered on vTM, the vTM can be configured to support the use of one or more applications. This can be achieved either by manually configuring the vTM using its GUI (see the *Virtual Traffic Manager* documentation) or by using *application templates*.

- [“Overview of Application Templates and Template Instances” on page 207.](#)
- [“Adding an Application Template to Services Director” on page 210.](#)
- [“Creating and Applying a Template Instance” on page 212.](#)
- [“Removing a vTM Application By Deleting a Template Instance” on page 218.](#)

Note: Application templates are only available to customers whose license includes the *Enterprise* Feature Tier.

Overview of Application Templates and Template Instances

A default configuration of resources and settings for a vTM application can be stored as an *application template* and uploaded into Services Director.

Note: Application templates are only available to customers whose license includes the *Enterprise* Feature Tier.

To use an application template on a registered vTM, Services Director creates a template instance from the application template. The exposed properties for the template instance are then previewed by the user, who can change any of the properties if required. These changes finalize the template instance.

The vTM-specific template instance is stored on the Services Director, and then applied automatically to the vTM cluster to create all required resources and settings for the application on all vTMs in the cluster.

Note: An application template can be used multiple times on a single vTM to create the resources and settings required for additional instances of the same application.

Example: Web Server Application Template

Services Director is supplied with an application template for a web server, which contains all required default information for a vTM-based web server application.

For this application template, a web server application requires:

- A port to receive incoming requests on the front-end IP address of the (clustered) vTM.
- Two back-end server pools to process the requests.

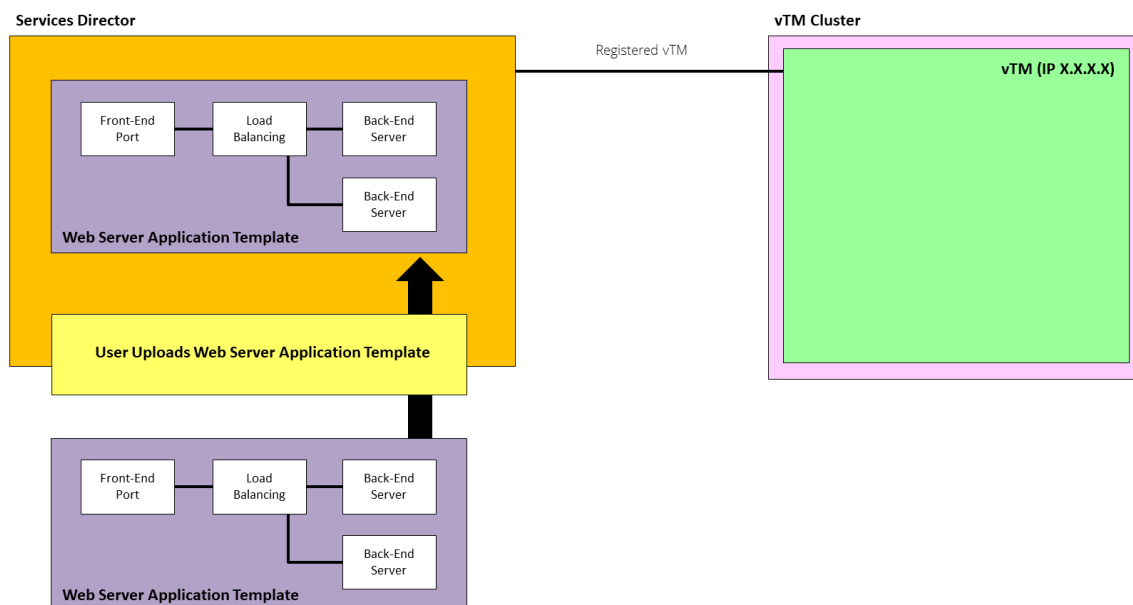
In this example:

- Two separate web servers are required on a single vTM.
- There is only one vTM in the cluster.

Note: Where multiple vTMs exist in the cluster, all vTMs in the cluster are configured for the application.


First, the web server template file (a .ZIP file) must be uploaded to the Services Director. For example:

FIGURE 167 Uploading an Application Template




After the application template is loaded into the Services Director, the user selects a vTM cluster to host the application. Services Director then creates a template instance from the application template. The exposed properties for the template instance are then previewed by the user, who can change any of the properties if required. These changes finalize the template instance. For example:

FIGURE 168 Creating a Template Instance



Services Director then applies the configuration from the template instance to the clustered vTM to create the first required web server application. For example:

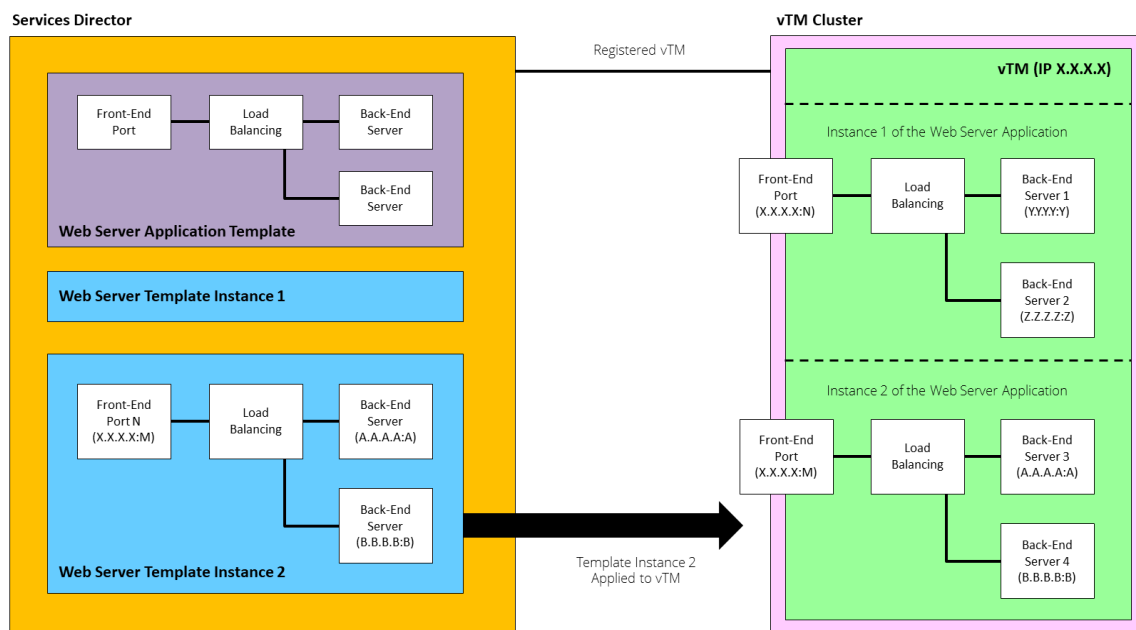
FIGURE 169 Creating a vTM Application from a Template Instance



Note: In this example, the back-end servers are implemented as vServers and Pools on the vTM.

For the second web server, the process can be repeated. A new template instance is always created. In this example, the template instance requires a different front-end port to the first template instance. For example:

FIGURE 170 Creating a Second Instance of a vTM Application



In this example:

- The Services Director has one application template and two template instances.
- The vTM cluster contains a single vTM.

Note: Where multiple vTMs exist in the cluster, all vTMs in the cluster are configured for the application.

- The vTM has two web server applications.

Adding an Application Template to Services Director

Services Director is supplied with application templates, which contain all required information to configure and create an instance of an application on a vTM. These are:

- Web server application template.
- SSL-based web server template.


Contact Pulse Secure to get these files.

To add an application template to Services Director:

1. Click the **Catalogs** menu, and then click **Application Templates**.

The **Application Templates** page appears. On its first use, this contains no entries. For example.


FIGURE 171 Empty Application Templates Page



2. Click the plus symbol above the application template table.

The **Import a Template** dialog box appears.

FIGURE 172 Import a Template



3. Select one of the following options:
 - **From URL.** Then, enter the URL for the application template.
 - **From File.** Then, click **Choose File** to locate the file.
4. Click **Apply**.

The application template is uploaded. After this completes, it is added to the **Application Templates** page. For example:

FIGURE 173 Application Templates Page

Application Templates

+ Add			
Name ▾	Version ▾	Description ▾	Date created ▾
▶ HTTP Service	1.0	A basic SSL decrypting web-service	2019-03-13 14:13:04

The uploaded application template is ready for use, see **“Creating and Applying a Template Instance” on page 212**.

5. (Optional) Expand the application template to see its full details. For example:

FIGURE 174 Application Template Details

Name ▾	Version ▾	Description ▾	Date created ▾
▼ HTTP Service	1.0	A basic HTTP web service	2019-05-16 12:55:56
Name	HTTP Service		
Version	1.0		
Description	A basic HTTP web service		
Author	www.pulsesecure.net		
Minimum vTM version required	18.2		
Date created	2019-05-16 12:55:56		

6. (Optional) Repeat steps 2 - 5 to add additional application templates.

After you have uploaded all required application templates, you can use them to create applications on vTMs in the estate of the Services Director, see **“Creating and Applying a Template Instance” on page 212**.

Creating and Applying a Template Instance

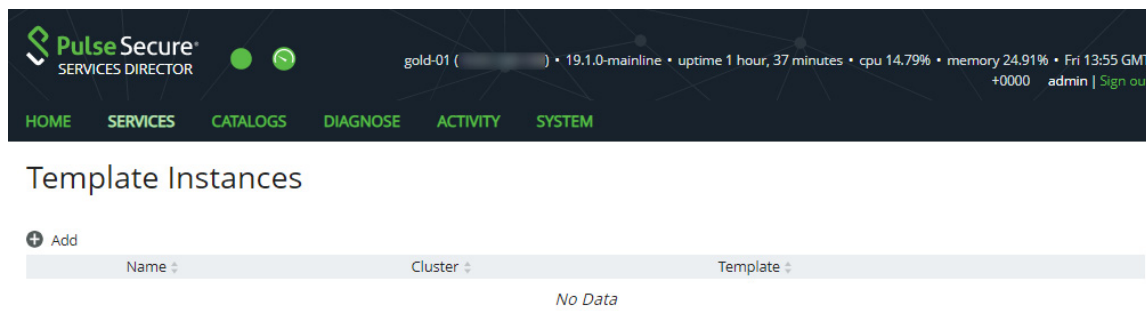
After you have uploaded one (or more) application templates to Services Director, you can create a template instance from an application template and apply the configuration to a vTM cluster.

To create and apply a template instance:

1. Click the **Services** menu, and then click **Application Templates: Template Instances**.

The **Template Instances** page appears. On its first use, this contains no entries. For example.

FIGURE 175 Empty Template Instances Page



2. Click the plus symbol above the template instances table.

The first page of the **Instantiate a template** wizard appears.

This page enables you to identify the required application template, and the required vTM cluster.

FIGURE 176 Instantiate a Template Wizard: Template/Cluster

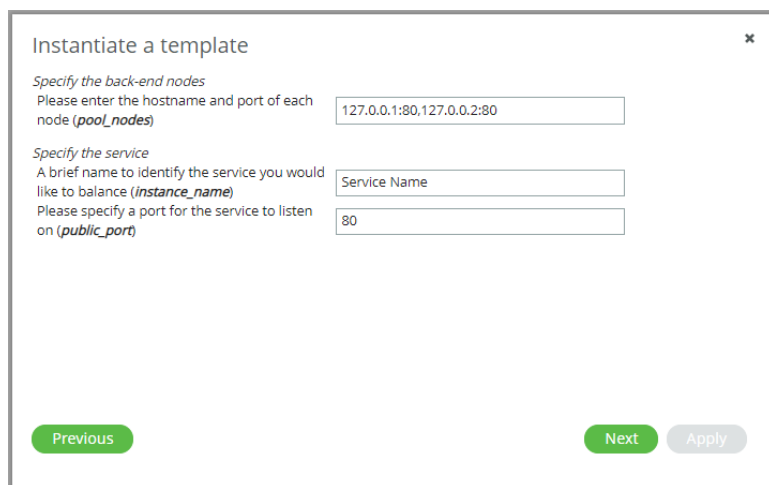
3. Select an application **Template** for the template instance.
4. Select a vTM **Cluster** for the template instance.
5. Enter a **Name** for the template instance.

6. Click **Next**.

The second page of the **Instantiate a template** wizard appears.

This page displays all properties that can be changed, as defined inside the template. Their default values for those properties are also displayed. For example:

FIGURE 177 Instantiate a Template Wizard: Properties



Instantiate a template

Specify the back-end nodes
Please enter the hostname and port of each node (*pool_nodes*)

127.0.0.1:80,127.0.0.2:80

Specify the service
A brief name to identify the service you would like to balance (*instance_name*)

Service Name

Please specify a port for the service to listen on (*public_port*)

80

Previous Next Apply

Note: This page of the wizard will vary between different application templates. For this reason, no property-specific instructions are given in this procedure.

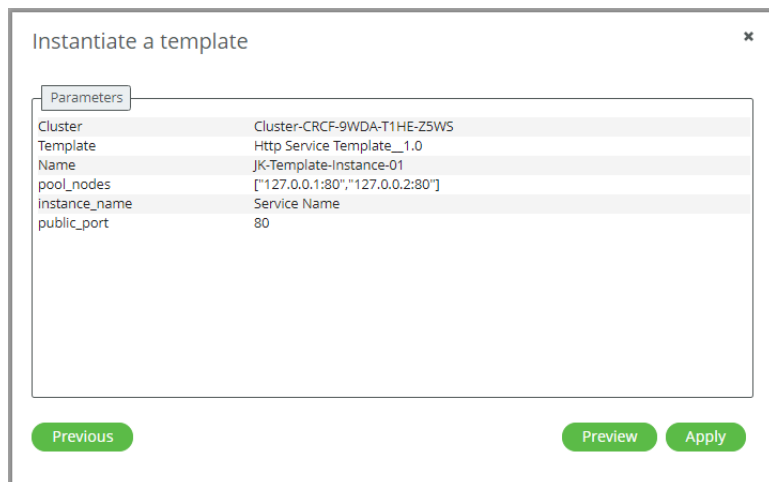
7. (Optional) Update any of the displayed values.

8. Click **Next**.

The third page of the **Instantiate a template** wizard appears.

This page summarizes the final values for each parameter. For example:

FIGURE 178 Instantiate a Template Wizard: Parameters



Instantiate a template

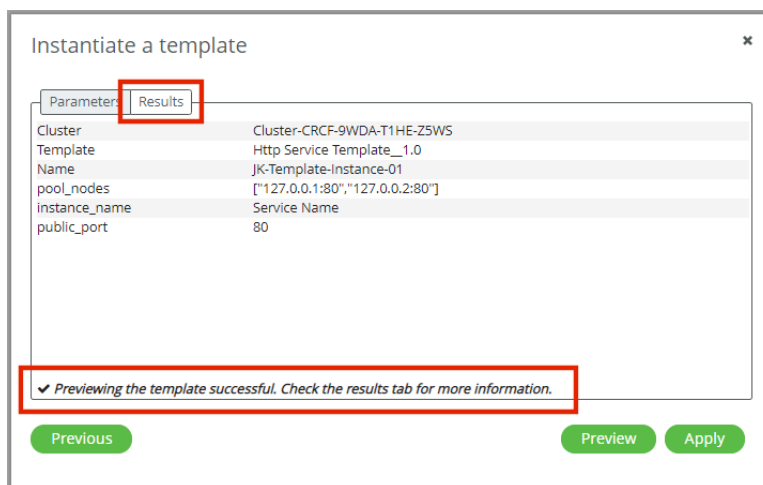
Parameters

Cluster	Cluster-CRCF-9WDA-T1HE-Z5WS
Template	Http Service Template_1.0
Name	JK-Template-Instance-01
pool_nodes	["127.0.0.1:80","127.0.0.2:80"]
instance_name	Service Name
public_port	80

Previous Preview Apply

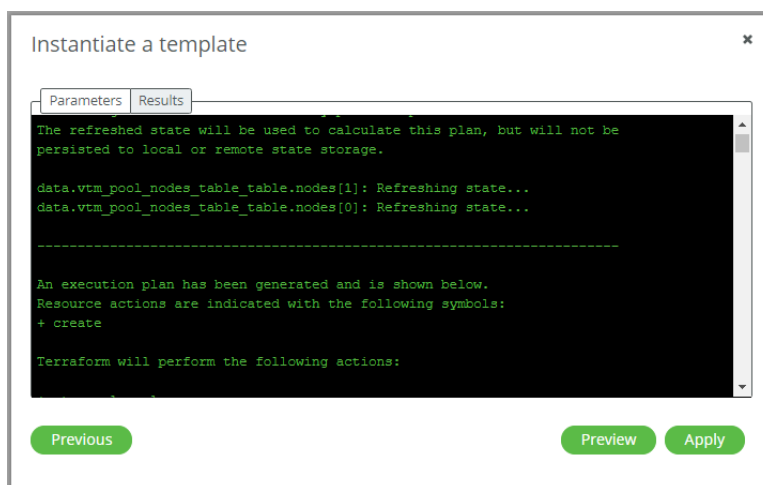
9. (Optional) Click **Preview** to test the template against the vTM settings for the cluster.
 - If the preview succeeds, the following message appears, and a **Results** tab is added.

FIGURE 179 Instantiate a Template Wizard: Review Success



(Optional) Click the **Results** tab to view the output of the preview operation. For example:

FIGURE 180 Instantiate a Template Wizard: Review Results



- If the preview fails, analyze the output in the **Results** tab and click **Previous** until you can change the properties for the template instance. Repeat as required.

10. Click **Apply**.

The template instance is created and the configuration is applied to all vTMs in the cluster.

11. Click the **Services** menu, and then click **Application Templates: Template Instances**.

The new template instance appears on the **Template Instances** page. For example:

FIGURE 181 Instantiate a Template Wizard: Parameters

Template Instances

+ Add		
Name	Cluster	Template
▶ HTTP-Service-01	Cluster-CRCF-9WDA-T1HE-Z5WS	HTTP Service

12. (Optional) To view details for the template instance, expand its entry.

FIGURE 182 Instantiate a Template Wizard: Parameters

Name	Cluster	Template
▼ HTTP-Service-01	Cluster-CRCF-9WDA-T1HE-Z5WS	HTTP Service

Cluster: Cluster-CRCF-9WDA-T1HE-Z5WS

Template: HTTP Service

Name:

Parameters:

- instance_name: "Service Name"
- public_port: 80
- nodes_list: ["127.0.0.1:80","127.0.0.2:80"]

13. (Optional) To confirm the application has been created correctly, log into a vTM in the cluster after a few minutes. For example, for a web server application:

- The **Services** summary on the **Home** page shows new vServers and pools.

FIGURE 183 Instantiate a Template Wizard: Confirming Services

The screenshot displays the Pulse Secure Services Director - Enterprise 19.1a2 interface. The top navigation bar includes the Pulse Secure logo, the product name, and a user '(admin/admin) Logout'. Below the navigation bar, a status bar shows 'Cluster: OK' and '0 b/s'. The main content area is divided into sections: 'Traffic Managers' showing two IP addresses (10.62.x.x) with status icons, and 'Services' showing two items: 'Web Server_virtual_server' (HTTP (80)) with a 'Running' status and a 'Web Server_pool' (Default Pool) with a status icon. A red rectangle highlights the 'Services' section.

- Click the virtual server **Service** to view the **Virtual Servers** tab. This tab shows the vServer properties specified in the template instance wizard.

FIGURE 184 Instantiate a Template Wizard: Confirming vServers

The screenshot displays the Pulse Secure Services Director web interface. The top navigation bar includes the Pulse Secure logo, the product name 'Virtual Traffic Manager Appliance Services Director - Enterprise 19.1a2', a user status bar with '(admin/admin) Logout', a cluster status 'Cluster: OK', and a bandwidth indicator '0 b/s'. Below the navigation bar, a breadcrumb trail shows 'Configuring: Traffic IP Groups > Virtual Servers > Web Server_virtual_server'. The main content area is titled 'Virtual Server: Web Server_virtual_server (HTTP, port 80)' and includes an 'Unfold All / Fold All' button. Under the heading 'Virtual Servers', it lists 'Pools used by this virtual server:' and shows a 'Web Server_pool' with a 'Default' status. The 'Last Modified' timestamp is '6 May 2019 10:04'. A 'Basic Settings' section is expanded, showing the following configuration: Name: 'Web Server_virtual_server', Enabled: 'Yes', Internal Protocol: 'HTTP', Port: '80', Default Traffic Pool: 'Web Server_pool', and Listening on: 'All IP addresses'. A 'Notes' field is present but empty. At the bottom of the settings panel are 'Update' and 'View traffic on World Map' buttons.

Pulse Secure® Virtual Traffic Manager Appliance Services Director - Enterprise 19.1a2 (admin/admin) Logout Cluster: OK 0 b/s

Configuring: Traffic IP Groups > Virtual Servers > Web Server_virtual_server Pools Config Summary

Virtual Servers

Virtual Server: Web Server_virtual_server (HTTP, port 80) Unfold All / Fold All

Pools used by this virtual server:

Web Server_pool Default

Last Modified: 6 May 2019 10:04

▼ Basic Settings

The basic settings specify the internal virtual server protocol that is used for traffic inspection, the port and IP addresses the virtual server listens on along with the default pool for handling traffic.

Name: Web Server_virtual_server

Enabled: ☒ Yes ☐ No

Internal Protocol: HTTP

Port: 80

Default Traffic Pool: Web Server_pool

Listening on: ☒ All IP addresses ☐ Traffic IP Groups ... ☐ Domain names and IP addresses ...

Notes:

Update View traffic on World Map

- The **Pools** tab shows the pool properties specified in the template instance wizard.

FIGURE 185 Instantiate a Template Wizard: Confirming Pools

Pulse Secure Virtual Traffic Manager Appliance Services Director - Enterprise 19.1a2 (admin/admin) Logout Cluster: OK 0 b/s

Configuring: Traffic IP Groups Virtual Servers **Pools > Web Server_pool** Config Summary

Pools Pool: Web Server_pool (HTTP) Unfold All / Fold All

Virtual servers that use this pool: Web Server virtual server Pools that use this pool as their failure pool: none

Last Modified: 6 May 2019 10:04

▼ **Basic Settings**

The basic settings specify the nodes to which the pool is balancing traffic.

Name: Web Server_pool

Node	State	Delete
127.0.0.1:80	Active	<input type="checkbox"/>
127.0.0.2:80	Active	<input type="checkbox"/>

Nodes:

Add Node(s):

Failure Pool: None

Notes:

Update

The creation of an application from an application template is complete.

Editing a Template Instance

After a application has been created, you can edit its properties in the **Template Instances** page:

- Click the **Services** menu, and then click **Application Templates: Template Instances**.

The list of template instances appears on the **Template Instances** page. For example:

FIGURE 186 Instantiate a Template Wizard: Parameters

Template Instances

+ Add		
Name	Cluster	Template
▶ HTTP-Service-01	Cluster-CRCF-9WDA-T1HE-Z5W5	HTTP Service

- (Optional) To view details for a template instance, expand its entry.

FIGURE 187 Instantiate a Template Wizard: Parameters

Name	Cluster	Template
HTTP-Service-01	Cluster-CRCF-9WDA-T1HE-Z5WS	HTTP Service

Cluster	Cluster-CRCF-9WDA-T1HE-Z5WS
Template	HTTP Service
Name	<input type="text" value="HTTP-Service-01"/> <input type="button" value="Update"/>
Parameters	instance_name "Service Name" public_port 80 nodes_list ["127.0.0.1:80","127.0.0.2:80"] <input type="button" value="Edit Parameters"/>

- Click **Edit Parameters**.

The **Update template instance parameters** wizard appears. For example:

FIGURE 188 Instantiate a Template Wizard: Properties

Update template instance parameters

Specify the back-end nodes
Please enter the hostname and port of each node (*nodes_list*)

Specify the service
A brief name to identify the service you would like to balance (*instance_name*)
Please specify a port for the service to listen on (*public_port*)

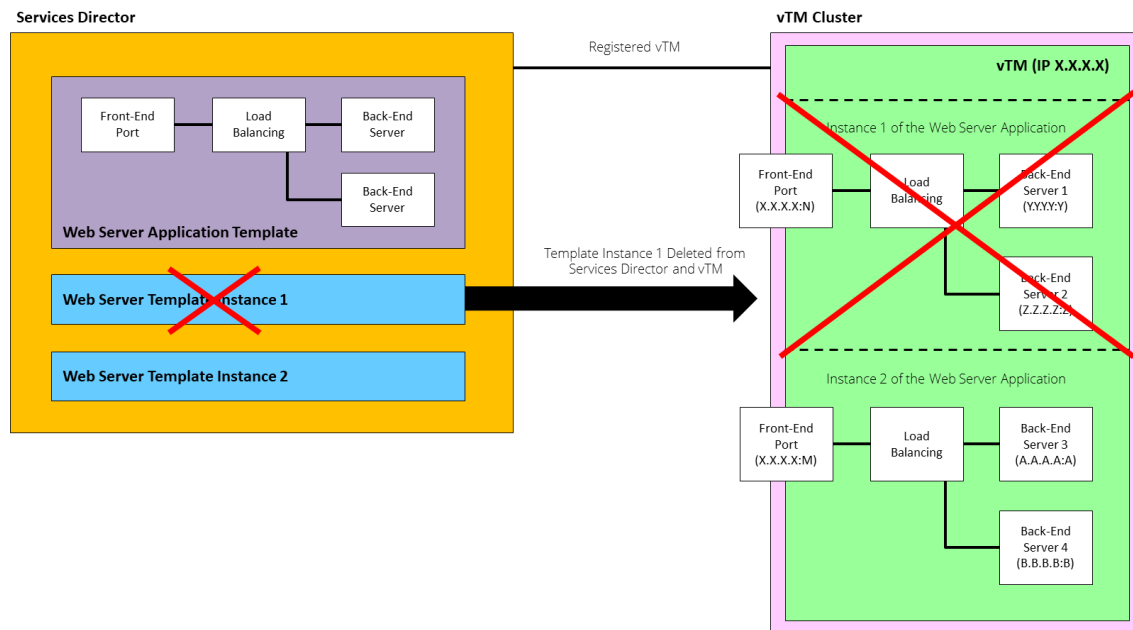
- Update the required values and continue with the wizard. This is the same as described in **“Creating and Applying a Template Instance” on page 212**.

Removing a vTM Application By Deleting a Template Instance

When you no longer require an application on a vTM that was configured from an application template, you can delete it. To do this, delete the matching application instance from the Services Director. Services Director automatically reconfigures the vTM, removing resources and resetting properties on the vTM so that the application is removed.

For example, where two web server applications exist on a vTM, if the first web server is no longer required, delete its matching template instance on the Services Director. Services Director automatically removes the resources and settings that were added for the first web server, but leaves the second web server intact. For example:

FIGURE 189 Removing a vTM Application By Deleting a Template Instance



To remove a vTM Application:

1. Click the **Services** menu, and then click **Application Templates: Template Instances**.

The list of template instances appears on the **Template Instances** page. For example:


FIGURE 190 Instantiate a Template Wizard: Parameters

Template Instances

+ Add		
Name ▾	Cluster ▾	Template ▾
▶ HTTP-Service-01	Cluster-CRCF-9WDA-T1HE-Z5WS	HTTP Service X

2. Ensure that no entries are expanded.
3. Hover the pointer over the template instance that you want to delete.
4. To the right of the template instance entry, click the **X** control. A confirmation control appears.

FIGURE 191 The Delete Confirmation Control




5. Click **Delete**.

The template instance is removed. Services Director automatically reconfigures the vTM, removing resources and resetting properties on the vTM so that the application is removed.

6. (Optional) To confirm the application has been deleted correctly, log into a vTM in the cluster after a few minutes and confirm that the removal is complete. For example, after the removal of a vTM's only web server application, ensure that the **Services** summary on the **Home** page shows the correct information. For example:

FIGURE 192 Instantiate a Template Wizard: Review Results



The removal of an application from a vTM is complete.

Relicensing Virtual Traffic Managers

Under a number of circumstances, you may need to relicense a vTM. For example:

- A Legacy FLA License is about to expire.
- The Service Endpoint Address of your Services Director changes. This affects vTMs that are licensed using either Universal FLA or Legacy FLA Licensing.
- A vTM is updated from version 10.0 (or earlier) to version 10.1 (or later). You can replace the Legacy FLA licensing with Universal FLA licensing.

Note: See [“Preparing to Relicense a Virtual Traffic Manager \(Legacy FLA to Universal FLA\)” on page 221](#) before starting this process.

- A new version of the Universal FLA License is released.
- The existing FLA License has been damaged in some way.

Note: If you are applying a new license to vTM that has no active REST API, you will need to add the Legacy FLA License to the vTM directly; this cannot be achieved through the Services Director.

Preparing to Relicense a Virtual Traffic Manager (Legacy FLA to Universal FLA)

You may have a vTM that you used on an earlier release of the Services Director, which is now at version 10.1 or later. You can change its current Legacy FLA Licensing to Universal FLA Licensing. Before you can do this, you must enable its REST API setting.

1. Click the **Services** menu, and then click **Services Director: vTM Instances**.

The **vTM Instances** page appears.

2. Locate the vTM's entry in the table of vTMs.
3. Click the arrow on the left side of the vTM's entry to show its details.
4. Under **vTM Management**, change **Rest API** to Enabled.
5. Click **Apply** to confirm the change.

You can then continue with the relicensing process.

Relicensing a Virtual Traffic Manager Instance

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.

- Click the **Catalogs** menu, and then click **Licensing: FLA Licenses**. The **FLA Licenses** page appears.

FIGURE 193 FLA Licenses Page

FLA Licenses

[+ Add License](#)

Universal Licenses

License Name ▾	Status ▾	Default ▾	Actions
▶ universal_v4	Active	Yes	Relicense

Legacy Licenses

License Name ▾	Status ▾	Default ▾	Actions
No Data			

- (Optional) Add any new flexible licenses. See [“Adding a Legacy FLA License to the Services Director” on page 144](#).
- Locate the license you wish to use.
This can be either a Universal FLA License or a Legacy FLA License.
- For this license, click **Relicense**.

The **Select Instances To Relicense** dialog box appears. This indicates the selected FLA License, and lists all current vTMs with an enabled REST API. For example:

FIGURE 194 FLA Licenses Page: vTM List

Select Instances to Relicense ✕

License name: (legacy)

☐ Select all

Relicense?	Name ▾	License Name ▾	Bandwidth ▾	Feature Pack ▾	Version ▾	Cluster ▾	Lifecycle ▾	Instance Health	Licensing Health
<input type="checkbox"/>	violet-02	universal_v3	100	STM-400_full		Cluster-RNPP-UIP9-RUA7-Q2JU	Active	OK	Licensed
<input type="checkbox"/>	violet-01	universal_v3	100	STM-400_full		Cluster-AC8L-ABCM-W5CR-ELSP	Active	OK	Licensed
<input type="checkbox"/>	sunshine-01	legacy_9.3	200	STM-400_full			Active	N/A	Licensed

[Relicense](#)

7. Select the required vTMs for the selected FLA License. For example:

FIGURE 195 Flexible Licenses Page: vTM Selections

Select Instances to Relicense

License name: (legacy)

☐ Select all

Relicense?	Name	License Name	Bandwidth	Feature Pack	Version	Cluster	Lifecycle	Instance Health	Licensing Health
<input type="checkbox"/>	violet-02	universal_v3	100	STM-400_full		Cluster-RNPP-UIP9-RUA7-Q2JU	Active	OK	Licensed
<input type="checkbox"/>	violet-01	universal_v3	100	STM-400_full		Cluster-AC8L-ABCM-W5CR-ELSP	Active	OK	Licensed
<input checked="" type="checkbox"/>	sunshine-01	legacy_9.3	200	STM-400_full			Active	N/A	Licensed

You may have a vTM that you used on an earlier release of the Services Director, which is now at version 10.1 or later. You can change its current Legacy FLA Licensing to Universal FLA Licensing. See [“Preparing to Relicense a Virtual Traffic Manager \(Legacy FLA to Universal FLA\)” on page 221](#).

8. Click **Relicense**. A confirmation dialog box appears.

FIGURE 196 FLA Licenses Page: vTM Confirmations

Select Instances to Relicense

License name: (legacy)

☐ Select all

Relicense?	Name	License Name	Bandwidth	Feature Pack	Version	Cluster	Lifecycle	Instance Health	Licensing Health
<input type="checkbox"/>	violet-02	universal_v3	100	STM-400_full		Cluster-RNPP-UIP9-RUA7-Q2JU	Active	OK	Licensed
<input type="checkbox"/>	violet-01	universal_v3	100	STM-400_full		Cluster-AC8L-ABCM-W5CR-ELSP	Active	OK	Licensed
<input checked="" type="checkbox"/>	sunshine-01	legacy_9.3	200	STM-400_full			Active	N/A	Licensed

Relicense Instances with new FLA license?

You have chosen to replace the FLA license on 1 vTM Instance.

Each indicated vTM will be transitioned from its original FLA license to the new FLA license.

Each relicensed Instance will encounter a brief interruption of service during the relicensing process.

Do you wish to continue?

9. Click **OK**. The relicensing process begins, and displays progress. There are two possible outcomes:
- The process completes successfully. For example:

FIGURE 197 FLA Licenses Page: vTM Relicensing Succeeds

Select Instances to Relicense

License name: (legacy)

☐ Select all

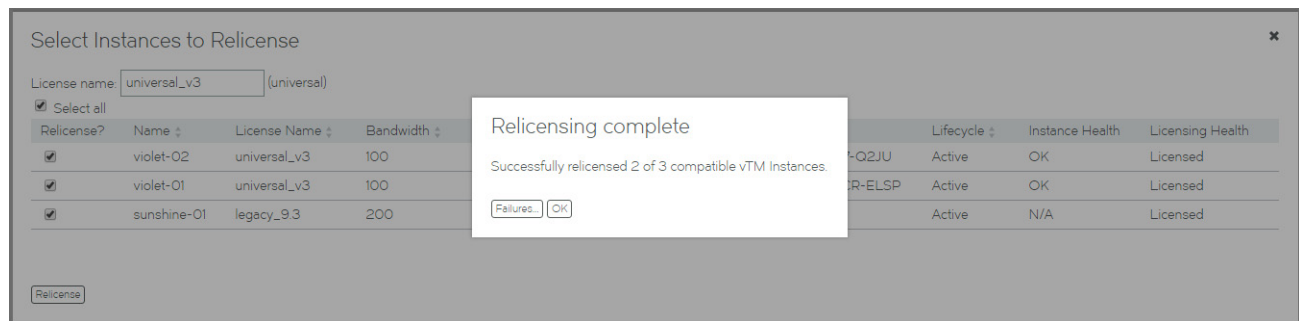
Relicense?	Name	License Name	Bandwidth	Feature Pack	Version	Cluster	Lifecycle	Instance Health	Licensing Health
<input type="checkbox"/>	violet-02	universal_v3	100	STM-400_full		Cluster-RNPP-UIP9-RUA7-Q2JU	Active	OK	Licensed
<input type="checkbox"/>	violet-01	universal_v3	100	STM-400_full		Cluster-AC8L-ABCM-W5CR-ELSP	Active	OK	Licensed
<input checked="" type="checkbox"/>	sunshine-01	legacy_9.3	200	STM-400_full			Active	N/A	Licensed

Relicensing complete

Successfully relicensed 1 of 1 compatible vTM Instances.

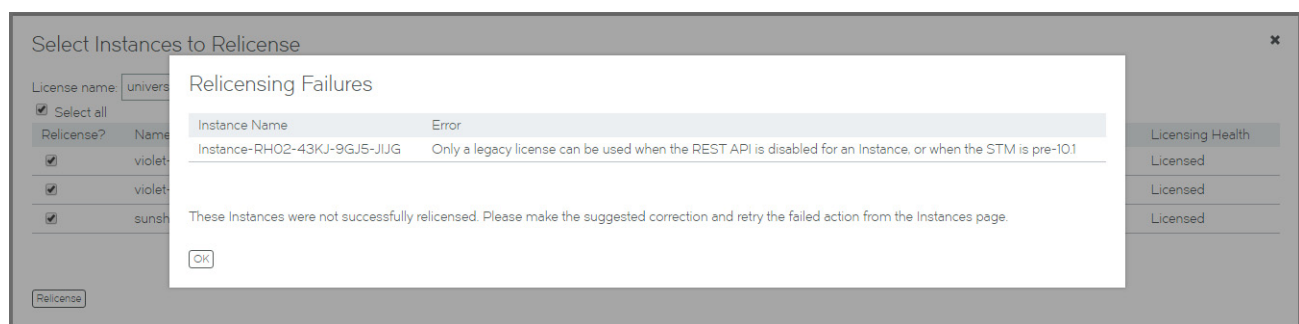
- The process completes, but is only partially successful. Using a different example:

FIGURE 198 FLA Licenses Page: vTM Relicensing Partial Completion



Click **Failures** to list the vTMs that could not be relicensed. For example:

FIGURE 199 FLA Licenses Page: vTM Relicensing Failures



You may need to investigate the licensing of these vTMs further.

10. Click **OK** to finish this process.

Processing Virtual Traffic Manager Metering Discrepancy Warnings

The accurate billing for Cloud Service Provider customers relies on:

- Accurate record-keeping for registered vTMs.
- Availability of metering information from each vTM.

The Services Director monitors the operation of each vTM to detect scenarios that may give rise to billing discrepancies. For example:

- A vTM was registered with the Services Director, but then decommissioned later without marking the vTM as *Deleted*. In this case, the decommissioned vTM will still be being charged on an uptime basis. This will result in over-accounting of uptime and a larger CSP bill than should have been charged.
- A vTM was registered with the Services Director, but the Services Director has been unable to retrieve metered throughput metrics from the vTM using its REST API or SNMP. In this case, the vTM will not have been charged for throughput at all. This is likely to result in under-metering and a smaller CSP bill than should have been charged.

Where no metering discrepancies are detected, the Services Director VA displays a green metering symbol in the header:

FIGURE 200 No Metering Discrepancies Detected



Where metering discrepancies are detected, the Services Director VA displays an orange metering warning symbol in the header:

FIGURE 201 Metering Discrepancy Warning



You can then inspect any metering warnings in the Services Director VA and resolve them. See [“Understanding Metering Discrepancy Warnings” on page 225](#).

Note: Monitoring that gives rise to metering alerts and notifications is enabled by default. You can change this setting if required from the **System > General Settings** page, see [“Updating Metering Alerts and Notifications Settings” on page 118](#).

Understanding Metering Discrepancy Warnings

Virtual Traffic Manager metering discrepancy warnings are displayed as a table in the **Metering Warnings** page.

To access this page, click the metering warning symbol in the header, see [“Processing Virtual Traffic Manager Metering Discrepancy Warnings” on page 224](#).

Alternatively, click the **Diagnose** menu and then click **Metering Warnings**.

In the **Metering Warnings** page, each line of the metering warnings table shows a potential billing discrepancy for a vTM. This includes:

- Timestamps for metering, licensing and monitoring.
- A summary reason for its inclusion.
- A potential solution, and the controls to access the solution.

For example:

FIGURE 202 Metering Warnings Page

Metering Warnings

NOTE

When connectivity to an instance is fixed, it will take up to 1 hour and 1 minute for the corresponding warning to disappear.

Name	Last Licensed	Last Monitored	Last Metered	Reason	Resolution	Shortcuts
cerise-01	2016-06-14 12:53:27	2016-06-14 12:54:35	2016-06-14 12:00:00	Possible uptime over-accounting	Mark instance as deleted if no longer in use	Delete
cerise-02	2016-06-16 12:55:00	2016-06-14 12:54:08	2016-06-14 12:00:00	Possible under-accounting	Enable REST or SNMP connectivity for this instance	Check connectivity Instance Settings
sienna-01	2016-06-14 12:53:39	2016-06-14 12:54:26	2016-06-14 12:00:00	Possible uptime over-accounting	Mark instance as deleted if no longer in use	Delete

In this example:

- There are two vTMs that are flagged as potentially being *over-billed*.

If a vTM is no longer in use, it is likely that it has not requested FLA licensing for over 24 hours, and cannot be contacted using REST API or SNMP. In this case, you can delete it to prevent over-billing for uptime. See [“Processing Potentially Over-Accounted Virtual Traffic Managers” on page 226](#).

- There is a vTM that is flagged as potentially being *under-billed*.

It is likely that this vTM is still requesting FLA licensing, but is uncontactable using REST API or SNMP. If you enable the REST API or SNMP for this vTM, this will re-enable metering and prevent under-billing for its use. See [“Processing Potentially Under-Accounted Virtual Traffic Managers” on page 227](#).

Note: Once these situations are resolved, the warnings and the warning symbol remain in place until the Services Director re-evaluates them. This may take up to one hour and one minute, and cannot be triggered from the interface.

Processing Potentially Over-Accounted Virtual Traffic Managers

If you are no longer using a vTM, but have not yet deleted it from the estate of the Services Director VA, you may see a metering discrepancy warning. This warning indicates that there is a possibility of the billing for the vTM being over-accounted. You can resolve this by deleting the vTM from the estate of the Services Director VA.

- In the header for the Services Director VA, click the metering warning symbol.

FIGURE 203 FLA Licenses Page



Alternatively, click the **Diagnose** menu and then click **Metering Warnings**.

The **Metering Warnings** page appears. This displays a table, with an entry for each vTM for which there is a metering discrepancy warning (see [“Understanding Metering Discrepancy Warnings” on page 225](#)).

2. Locate the entry for the required vTM.
3. Examine the registered details for the vTM.

To do this, visit the **vTM Instances** page and/or examine the user interface of the vTM itself.

4. If you decide to delete the vTM, click **Delete** in the **Shortcuts** column.

The entry is marked as *Deleted* in the **Shortcuts** column. Then, after a short time, the entry is removed from the table.

Processing Potentially Under-Accounted Virtual Traffic Managers

The Services Director VA uses the REST API to collect metering information. If the REST API is not enabled, SNMP is then attempted if your configuration supports it. If you are using a vTM without either its REST API or SNMP active, you may see a metering discrepancy warning. This warning indicates that there is a possibility of the billing for the vTM being under-accounted. You can resolve this by enabling the REST API or SNMP for the vTM.

1. In the header for the Services Director VA, click the metering warning symbol.

FIGURE 204 FLA Licenses Page



Alternatively, click the **Diagnose** menu and then click **Metering Warnings**.

The **Metering Warnings** page appears. This displays a table, with an entry for each vTM for which there is a metering discrepancy warning (see [“Understanding Metering Discrepancy Warnings” on page 225](#)).

2. Locate the entry for the required vTM.
3. Click **Instance Setting** for the entry.

The **vTM Instances** page appears.

4. In the table of vTMs on the **vTM Instances** page, expand the vTM to show its detailed view.
5. Check the **REST API**, REST Address and SNMP Address settings in the detailed view.
6. If the **REST API** is Disabled, the REST API has been disabled from the Services Director VA. Set this to Enabled and **Apply** the change.

Note: Once the REST API for the vTM shows as Enabled on the **Metering Warnings** page, it is not guaranteed that the REST API is enabled on the vTM itself. You must continue with this procedure to the end to ensure its operation.

7. In the detail view for the vTM, click **Please click for more details**.

You are redirected to the vTM's login page.

8. If you want to use the REST API to gather metering information, enable it on the vTM. Refer to the Virtual Traffic Manager documentation for details.
9. If you want to use SNMP to gather metering information, enable it on the vTM. Refer to the Virtual Traffic Manager documentation for details.
10. Return to the **Metering Warnings** page on the Services Director VA.
11. For the required vTM, click **Check connectivity**.

The connectivity between the Services Director VA and the vTM is tested. If this test succeeds, Check successful appears.

Note: The vTM entry is not removed from the table immediately. This can take up to one hour and one minute.

Working with Virtual Traffic Manager Clusters

- [Overview: Working with Virtual Traffic Manager Clusters](#) 229
- [Understanding Virtual Traffic Manager Cluster Details](#) 230
- [Creating a Virtual Traffic Manager Cluster](#) 233
- [Updating a Virtual Traffic Manager Cluster](#) 234
- [Working with vTM Cluster Backups](#) 237
- [Moving a vTM Between Clusters](#) 257
- [Deleting an Empty Virtual Traffic Manager Cluster](#) 258

Overview: Working with Virtual Traffic Manager Clusters

The **vTM Cluster** page displays a list of all Virtual Traffic Manager (vTM) clusters known to the Services Director VA.

The **vTM Cluster** page also enables you to:

- Assign an analytics profile to the cluster, which enables vTM analytics on all vTMs in the cluster. See [“Configuring vTM Analytics on the Services Director” on page 281](#).
- Assign a backup schedule to each cluster.
- Inspect the details of the cluster backups taken.

FIGURE 205 vTM Clusters Page

vTM Clusters

+ Add								
Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
▶ Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	N/A	shed-hourly-01	2016-07-03 08:30:00	Backup Now		
▶ Cerise-Cluster	Discovered	✓	N/A	N/A		Backup Now		
▶ Cluster-RNPP-UIP9-RUA7-Q2JU	Discovered	✓	N/A	N/A		Backup Now		
▶ Violet-Cluster	User Created		N/A	N/A		Backup Now		

There are two types of clusters used by the Services Director VA:

- *Discovered* - this is a cluster present on one or more externally-deployed vTMs. When an externally-deployed vTM is registered, a cluster name is displayed automatically.

Note: Registering a clustered vTM does not register other vTMs in the cluster, nor does it license them; you must independently register and license each node in a cluster.

Note: You cannot create a Discovered cluster from the **vTM Clusters** page.

Note: Services Director's awareness of Discovered clusters is limited to vTMs at version 10.2 or later with an enabled REST API.

- *User Created* - this is a cluster that you create manually on the **vTM Clusters** page. This cluster type can *only* be used for vTMs that you deploy from the Services Director VA. Refer to the *Pulse Services Director Advanced User Guide* for details.

You can rename a cluster of either type from the **vTM Clusters** page, see [“Updating a Virtual Traffic Manager Cluster” on page 234](#).

Services Director supports backup and restore for cluster configurations, see [“Working with vTM Cluster Backups” on page 237](#).

Understanding Virtual Traffic Manager Cluster Details

The **vTM Cluster** page displays a table of clusters known to the Services Director VA.

FIGURE 206 vTM Cluster Page

vTM Clusters

+ Add									
Cluster Name ↑	Type ↓	In Use ↓	Analytics Profile ↓	Backup Schedule ↓	Next Backup Time ↓	Action	Last Action ↓	Last Action Status	
▶ Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	N/A	shed-hourly-01	2016-07-03 08:30:00	Backup Now			
▶ Cerise-Cluster	Discovered	✓	N/A	N/A		Backup Now			
▶ Cluster-RNPP-UIP9-RUA7-Q2JU	Discovered	✓	N/A	N/A		Backup Now			
▶ Violet-Cluster	User Created		N/A	N/A		Backup Now			

Each entry in the table of clusters on the **vTM Clusters** page shows basic details for each cluster, and provides controls for backup operations where supported by the cluster.

Name	Description
Cluster Name	<p>The unique name of the cluster.</p> <p>If required, you can rename a cluster. See “Updating a Virtual Traffic Manager Cluster” on page 234.</p>
Type	<p>There are two cluster types used by the Services Director:</p> <ul style="list-style-type: none"> • <i>Discovered</i> - this is a cluster present on one or more externally-deployed vTMs. When an externally-deployed vTM is registered (version 10.2 or later with an active REST API), a cluster name is displayed automatically. <p>Note: Registering a clustered vTM does not register other vTMs in the cluster, nor does it license them; you must independently register and license each node in a cluster.</p> <p>Note: You cannot create a Discovered cluster from the vTM Clusters page.</p> <p>Note: Services Director's awareness of Discovered clusters is limited to vTMs at version 10.2 or later with an enabled REST API.</p> <ul style="list-style-type: none"> • <i>User Created</i> - this is a cluster that you create manually on the vTM Clusters page. This cluster type can <i>only</i> be used for vTMs that are deployed from the Services Director VA. Refer to the <i>Pulse Services Director Advanced User Guide</i> for details.

Name	Description
In Use	This indicates whether any vTMs are currently in the cluster.
Analytics Profile	(Optional) The assigned analytics profile for the cluster. See “Configuring vTM Analytics on the Services Director” on page 281 .
Backup Schedule	<p>(Optional) The selected schedule for the cluster backup. The configured number of backups for this cluster and the most recent backups are displayed in the detail view for the cluster. See “Creating a Cluster Backup Schedule” on page 239.</p> <p>Note: Where no Backup Schedule is selected, this property is displayed as <i>N/A</i>.</p> <p>Note: This column is only supported on vTMs at version 11.0 and later.</p> <p>Note: The number of backups for this cluster is visible in the detail view for the cluster.</p>
Next Backup Time	<p>The time of the next scheduled cluster backup.</p> <p>Note: Where no Backup Schedule is selected, this property is blank.</p> <p>Note: This column is only supported on vTMs at version 11.0 and later. This column is blank for all other vTMs.</p>
Action	<p>This column displays buttons that activate (or report on) supported cluster backup activities. This includes:</p> <ul style="list-style-type: none"> • Backup Now. When clicked, a backup is performed immediately. • Retry. This appears after a user-triggered Backup Now action fails. When clicked, the Backup Now action is re-attempted. See “Retrying An Immediate Backup After a Failure” on page 247. • Clear Failed Action. This appears after a user-triggered Backup Now action fails. When clicked, both the named Last Action and the Failed Last Action Status are removed. See “Retrying An Immediate Backup After a Failure” on page 247. <p>Note: This column is only supported on vTMs at version 11.0 and later. Where the vTM does not support backups, the Backup Now button is displayed but remains unavailable.</p>
Last Action	<p>The most recent manually-performed Action for a cluster backup (see above). This can be:</p> <ul style="list-style-type: none"> • Backup Now. This appears after a Backup Now action is attempted (see above). • Restore. This appears after a restore operation is attempted for a listed cluster backup. See “Restoring a Backup to a Cluster” on page 250. • Upload. This appears after an upload operation is attempted for a listed cluster backup. See “Uploading a Cluster Backup to a Virtual Traffic Manager” on page 253. <p>The result of the displayed action is shown in the Last Action Status column (see below). Scheduled backups are not included in this column.</p> <p>Note: This column is only supported on vTMs at version 11.0 and later. This column is blank for all other vTMs.</p>
Last Action Status	<p>The outcome of the Last Action operation (see above). This is blank, <i>In Progress</i> (blue), <i>Complete</i> or <i>Failed</i> (red).</p> <p>The results of scheduled backups are not included in this column.</p> <p>Note: A failed flag can be cleared from the Action column (see above).</p> <p>Note: This column is only supported on vTMs at version 11.0 and later. This column is blank for all other vTMs.</p>

To view the full details for a cluster, expand the required cluster. This includes:

- a **Cluster Name** that you can update, see [“Updating a Virtual Traffic Manager Cluster” on page 234](#).
- an **Owner** for the cluster.
- the **Analytics Profile** for the cluster. See [“Configuring vTM Analytics on the Services Director” on page 281](#).
- the **Backup Schedule** and **Number of Backups** that define the backup schedule for the cluster, where one is used. See [“Overview: vTM Cluster Backups” on page 237](#).

For example, when no cluster backup is in use:

FIGURE 207 Cluster Detailed View: No Backups Present

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	N/A	sched-hourly-01	2016-07-03 08:30:00	Backup Now		
Cerise-Cluster	Discovered	✓	N/A	N/A		Backup Now		

Cluster Name:

Owner:

Analytics Profile:

Backup Schedule:

Number of Backups:

Apply

Revert

Backup Name	Description	Date	Retain	Actions
There are no backups currently available for this cluster				

Where a cluster was created for a cloud-based vTM, an additional field containing an AWS user data block is included.

FIGURE 208 Cluster Detailed View: Cluster for Cloud-Based vTMs

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
▼ AWS-cluster-01	Discovered	✓	N/A	N/A		Backup Now		

Cluster Name:

Owner:

Analytics Profile:

Backup Schedule:

Number of Backups:

AWS User Data for Instances to join this Cluster:

```
Y2x1c3Rlc19ob3N0PTEwLjguMi4xMTUKY2x1c3Rlc19maW
5nzX3wcm1udd0yRTowQzozOTpBNToxQj05MjpeNzozQT00
NTPCODpERDoxRD0z03MjpbOT0yQTpCQTpCNjpfFQTow0D040Q

```

Copy to clipboard

Apply

Revert

This AWS user data text block is required when you create additional cloud-based vTM cluster members, see [“Creating the Second vTM in a Cluster” on page 190](#).

Use **Copy to clipboard** before performing this task.

Where a backup schedule for the cluster is in use, a list of backups is included. For example:

FIGURE 209 Cluster Detailed View: Backups Present

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-hourly-01	2016-07-03 08:30:00	Backup Now		

Cluster Name:
Owner:
Analytics Profile:
Backup Schedule:
Number of Backups:

Backup Name	Description	Date	Retain	Actions
▶ Backup-QB4D-QTRH-N6VE-LTLZ	Cluster-AQJE-R4HV-QYRI-9F4O#74	2016-07-03 06:30		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-NXIJ-BPOY-U4F6-OS28	Cluster-AQJE-R4HV-QYRI-9F4O#75	2016-07-03 07:30		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-MPFV-K23G-H3FT-OVWE	Cluster-AQJE-R4HV-QYRI-9F4O#72	2016-07-03 04:30		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-YFAP-9YU9-T75R-3ZAV	Cluster-AQJE-R4HV-QYRI-9F4O#71	2016-07-03 03:30		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-ZUOZ-TTF3-NKEZ-FJDR	Cluster-AQJE-R4HV-QYRI-9F4O#73	2016-07-03 05:30		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>

To make use of any listed backups, see [“Working with vTM Cluster Backups” on page 237](#).

Creating a Virtual Traffic Manager Cluster

You can create a *User Created* vTM cluster from the **vTM Clusters** page.

Note: You cannot create a *Discovered* cluster using the Services Director.

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **Services Director: vTM Clusters**. The **vTM Clusters** page appears.
4. Click the plus symbol above the vTM cluster table.

The **Add vTM Cluster** dialog box appears.

FIGURE 210 Adding a vTM Cluster

5. Specify the following:

- **Cluster Name** - specify the unique name for the cluster.
- **Owner** - select an owner for the cluster.
Note: If there are no owner entries, see [“Adding an Owner to the Services Director” on page 141](#).
- **Analytics Profile** - (Optional) Specify an analytics profile for the cluster. See [“Configuring vTM Analytics on the Services Director” on page 281](#).
- **Backup Schedule** - (Optional) Select an existing backup schedule. If you want to create a new schedule, click **Add new schedule**. When you do this, this page is replaced by the **Instances Backup Schedule** page. See [“Creating a Cluster Backup Schedule” on page 239](#).

6. Click **Add**.

The *User Created* cluster is added to the table of clusters.

Updating a Virtual Traffic Manager Cluster

You can update a vTM cluster from the **vTM Clusters** page.

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **Services Director: vTM Clusters**. The **vTM Clusters** page appears.

- Expand the required cluster.

FIGURE 211 Updating a vTM Cluster

vTM Clusters

Add

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-RNPP-UIP9-RUA7-Q2JU	Discovered	✓	N/A	N/A		Backup Now		
Cluster-P2WL-IV2B-V8S6-COIY	Discovered	✓	N/A	N/A		Backup Now		

Cluster Name:

Owner:

Analytics Profile:

Backup Schedule:

Number of Backups:

Backup Name	Description	Date	Retain	Actions
There are no backups currently available for this cluster				

- Update the **Cluster Name**. For example:

FIGURE 212 Specifying a New Name For a vTM Cluster

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-RNPP-UIP9-RUA7-Q2JU	Discovered	✓	N/A	N/A		Backup Now		
Cluster-P2WL-IV2B-V8S6-COIY	Discovered	✓	N/A	N/A		Backup Now		

Cluster Name:

Owner:

Analytics Profile:

Backup Schedule:

Number of Backups:

- (Optional) Select both a new **Backup Schedule** and a **Number of Backups**. See [“Working with vTM Cluster Backups” on page 237](#).

Note: The **Number of Backups** property is only used when there is a **Backup Schedule** selected.

- Click **Apply**. The cluster is updated.

FIGURE 213 vTM Cluster Page: Updated (Renamed) Cluster

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action
Cluster-RNPP-UIP9-RUA7-Q2JU	Discovered	✓	N/A	N/A		Backup Now	
Cerise-Cluster	Discovered	✓	N/A	N/A		Backup Now	
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	shed-hourly-01	2016-07-01 11:30:00	Backup Now	

To view updated **Backup Schedule** and **Number of Backups** settings, expand the cluster.

You can also confirm the name change from the **vTM Instances** page. For example:

FIGURE 214 Confirming an Updated vTM Cluster

vTM Instances

► Filters Filtering by Lifecycle, Instance Health, Licensing Health

⊕ Add

Show: 20 ▼ of 5 instances

Name ↕	License Name ↕	Bandwidth ↕	Feature Pack ↕	Version ↕	Cluster ↕	Instance Lifecycle ↕	Instance Health ↕	Licensing Health ↕
▶ sienna-01	universal_v4	120	STM-400_full	11.0	Cluster-AQJE-R4HV-QYRI-9F4O	Active	OK	Licensed
▶ sienna-02	universal_v4	120	STM-400_full	11.0	Cluster-AQJE-R4HV-QYRI-9F4O	Active	OK	Licensed
▶ violet-01	universal_v4	80	STM-400_full	10.3	Cluster-RNPP-UIP9-RUA7-Q2JU	Active	OK	Licensed
▶ cerise-01	universal_v4	100	STM-400_full	10.4	Cerise-Cluster	Active	OK	Licensed
▶ cerise-02	universal_v4	100	STM-400_full	10.4	Cerise-Cluster	Active	OK	Licensed

In this example, the *Cerise-Cluster* name is shown for both vTMs that are in the cluster.

Working with vTM Cluster Backups

All of the vTMs in a cluster share a cluster configuration. To ensure that the cluster configuration is preserved, you can schedule a regular cluster backup for each cluster. This preserves the cluster configuration only, and not the individual configuration of each vTM. This section includes the following topics:

- [“Overview: vTM Cluster Backups” on page 237](#)
- [“Creating a Cluster Backup Schedule” on page 239](#)
- [“Updating a Cluster Backup Schedule” on page 241](#)
- [“Adding a Backup Schedule to a Cluster” on page 242](#)
- [“Viewing Backups for a Cluster” on page 243](#)
- [“Updating Details for a Cluster Backup” on page 245](#)
- [“Performing an Immediate Backup for a Cluster” on page 246](#)
- [“Comparing Two Cluster Backups” on page 248](#)
- [“Restoring a Backup to a Cluster” on page 250](#)
- [“Uploading a Cluster Backup to a Virtual Traffic Manager” on page 253](#)
- [“Deleting a Cluster Backup” on page 256](#)

Note: The use of Cluster Backups is optional, and is only available to customers who license analytics features.

Note: Cluster Backups are not the same as Services Director backups. Services Director backups enable you to recover from a Services Director failure, see [“Recovering from a Services Director Failure” on page 401](#).

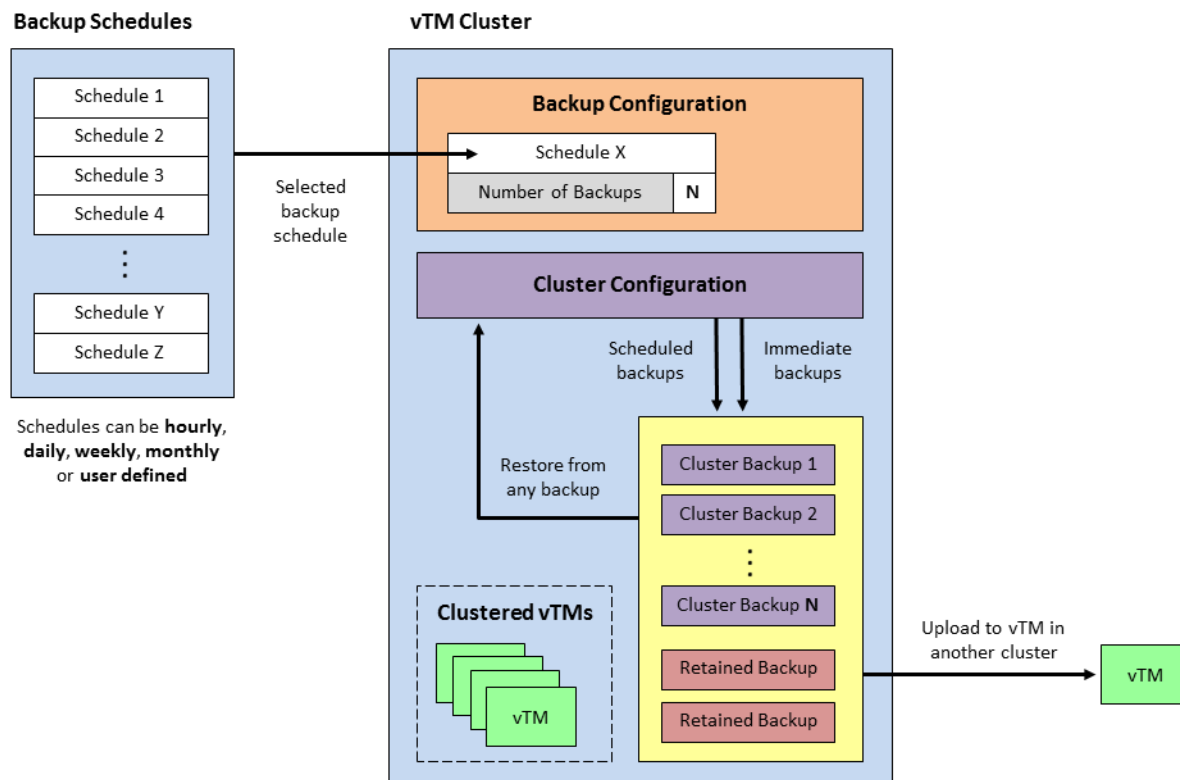
Overview: vTM Cluster Backups

A vTM cluster gathers vTMs together and operates them under a shared cluster configuration.

The configuration of the cluster can be backed up automatically on a regular basis according to a backup schedule.

The following provides an overview of automatic cluster backup operations.

FIGURE 215 Overview of Cluster Backups



Before you set up automatic backups for a cluster's configuration, you must create one or more backup schedules, see [“Creating a Cluster Backup Schedule” on page 239](#). Backup schedules define the frequency and times at which a backup will be taken. Each can be applied to one or more clusters.

Once you have backup schedules, you can configure the cluster to create backups automatically using a backup schedule. To do this, you select a backup schedule for the cluster, and indicate the number of backups that you want the cluster to store, see [“Adding a Backup Schedule to a Cluster” on page 242](#).

Once the cluster has an assigned cluster backup schedule, the cluster accumulates scheduled backups automatically. You can also manually request an immediate backup at any time. See [“Performing an Immediate Backup for a Cluster” on page 246](#).

Note: You can also request an immediate backup when there is no assigned backup schedule.

Once the maximum number of cluster backups is reached, older cluster backups are deleted automatically whenever newer cluster backups are created.

You can also choose to *retain* one or more backups if required, see [“Updating Details for a Cluster Backup” on page 245](#). Retained backups do not count towards the maximum number of backups for the cluster, and are not deleted automatically.

The cluster's configuration can be restored from an existing backup at any time, see [“Restoring a Backup to a Cluster” on page 250](#).

To support the selection of the correct cluster backup, you can compare any two cluster backups to identify the differences, see [“Comparing Two Cluster Backups” on page 248](#).

Also, you can upload a cluster backup to any vTM known to the Services Director, see [“Uploading a Cluster Backup to a Virtual Traffic Manager” on page 253](#). The uploaded configuration file is stored by the vTM, but not restored. This enables you to perform additional analysis and comparison using the vTM's graphical user interface.


Creating a Cluster Backup Schedule

A cluster backup schedule is a definition of when a cluster backup will be created. This includes general frequency (hourly, daily, weekly, monthly, and instant backups) and information to specify an exact backup time.

Defined schedules are displayed in the **vTM Backup Schedules** page. For example:

FIGURE 216 The vTM Backup Schedules Page

vTM Backup Schedules

 Add

Schedule Name ↕	Frequency ↕	Backup Time ↕	Details ↕
▶ sched-daily-01	Daily	10:10	Daily backup schedule
▶ sched-hourly-01	Hourly	N/A	Hourly backup schedule
▶ sched-monthly-01	Monthly	11:30	Monthly (11th) backup schedule
▶ sched-user-01	Every 12 Hours (Instant Backup)	14:16	12-hourly backup schedule
▶ sched-weekly-01	Weekly	10:10	Weekly backup schedule
▶ sched-weekly-02	Weekly	12:16	Weekly backup schedule (Friday)

Once you have created a schedule, it can be applied to any clusters that require the specified backup schedule, see [“Adding a Backup Schedule to a Cluster” on page 242](#).

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **Services Director: vTM Backup Schedules**. The **Instances Backup Schedule** page appears.
4. Click the plus sign above the table of backup schedules.

The **Add vTM Backup Schedule** dialog box appears.

FIGURE 217 Creating a Cluster Backup Schedule

5. Specify the required **Schedule Name** for the backup schedule.
6. (Optional) Enter a description for the backup schedule as its **Schedule Info**.

Note: This will be displayed as **Details** in the table of schedules.

7. Select the required **Frequency** for the backup schedule:
 - **Hourly** - this schedule will be performed once every hour. By default, this is on the hour. You can also choose to **Schedule At** 15, 30 and 45 minutes past the hour.
 - **Daily** - this schedule will be performed once per day. By default, this is at midnight. Alternatively, you can choose to **Schedule At** a specific time (hh:mm).
 - **Weekly** - this schedule will be performed once per week. By default, this is on Monday at midnight. Alternatively, you can choose to **Schedule On** the required day (Monday - Sunday) and **Schedule At** a specific time (hh:mm).
 - **Monthly** - this schedule will be performed once per month. By default, this is on Monday at midnight. Alternatively, you can choose to **Schedule On** the required day (typically, 1-28) and **Schedule At** a specific time (hh:mm).
 - **Instant Backup** - this schedule will be performed at a custom frequency. Instead of specifying an exact time, the first backup will be taken immediately when the schedule is applied to a cluster, and then at the defined **Schedule Every** frequency: every 15 minutes, hourly, every 12 hours, every week, every month).
8. Click **Add**.

The new schedule is added to the table of backup schedules.

Once you have created a schedule, it can be applied to any clusters that require the specified backup schedule, see [“Adding a Backup Schedule to a Cluster” on page 242](#).

Updating a Cluster Backup Schedule


Once a cluster backup schedule is created, you can change it at any time. The schedule can be renamed, and any of the schedule details can be changed.

Any cluster that uses the backup schedule will automatically make use of the revised updated schedule.

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **Services Director: vTM Backup Schedules**. The **vTM Backup Schedules** page appears.
4. Expand the required cluster backup schedule. For example:

FIGURE 218 Updating a Cluster Backup Schedule

vTM Backup Schedules

 Add

Schedule Name	Frequency	Backup Time	Details
▶ sched-daily-01	Daily	10:10	Daily backup schedule
▼ sched-hourly-01	Hourly	N/A	Hourly backup schedule

Schedule Name:

Schedule Info:

Frequency: ☒ Hourly ☐ Daily ☐ Weekly
☐ Monthly ☐ Instant backup

Schedule At: minutes past the hour

▶ sched-monthly-01	Monthly	11:30	Monthly (11th) backup schedule
▶ sched-user-01	Every 12 Hours (Instant Backup)	14:16	12-hourly backup schedule
▶ sched-weekly-01	Weekly	10:10	Weekly backup schedule
▶ sched-weekly-02	Weekly	12:16	Weekly backup schedule (Friday)

5. (Optional) Specify a new **Schedule Name** for the backup schedule.
6. (Optional) Enter a new description for the backup schedule as its **Schedule Info**.

Note: This will be displayed as **Details** in the table of schedules.

7. (Optional) Select a new **Frequency** for the backup schedule:
 - **Hourly** - this schedule will be performed once every hour. By default, this is on the hour. You can also choose to **Schedule At** 15, 30 and 45 minutes past the hour.
 - **Daily** - this schedule will be performed once per day. By default, this is at midnight. Alternatively, you can choose to **Schedule At** a specific time (hh:mm).

- **Weekly** - this schedule will be performed once per week. By default, this is on Monday at midnight. Alternatively, you can choose to **Schedule On** the required day (Monday - Sunday) and **Schedule At** a specific time (hh:mm).
- **Monthly** - this schedule will be performed once per month. By default, this is on Monday at midnight. Alternatively, you can choose to **Schedule On** the required day (typically, 1-28) and **Schedule At** a specific time (hh:mm).
- **Instant Backup** - this schedule will be performed at a custom frequency. Instead of specifying an exact time, the first backup will be taken immediately when the schedule is applied to a cluster, and then at the defined **Schedule Every** frequency: every 15 minutes, hourly, every 12 hours, every week, every month).

Note: If your **Schedule Name** and **Schedule Info** include references to the Frequency, remember to update these also.

8. Click **Apply**.

The schedule is updated in the table of backup schedules.

Note: Any cluster that uses the backup schedule will automatically make use of the revised updated schedule.

Adding a Backup Schedule to a Cluster

Once you have created a cluster backup schedule (see [“Creating a Cluster Backup Schedule” on page 239](#)), it can be applied to one or more clusters. This ensures that the required cluster backup schedule is performed for all of those clusters.

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **Services Director: vTM Clusters**. The **vTM Clusters** page appears.

- Expand the required cluster.

FIGURE 219 Cluster Without Backup Schedule

vTM Clusters

Add

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-RNPP-UIP9-RUA7-Q2JU	Discovered	✓	N/A	N/A		Backup Now		
Violet-Cluster	User Created		N/A	N/A		Backup Now		
Cerise-Cluster	Discovered	✓	N/A	N/A		Backup Now		
Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	N/A	N/A		Backup Now		

Cluster Name:

Owner:

Analytics Profile:

Backup Schedule:

Number of Backups:

Backup Name	Description	Date	Retain	Actions
There are no backups currently available for this cluster				

- Select the required **Backup Schedule**.
- Specify the required **Number of Backups**. The default is 5.

Note: Retained backups are not included in this number. See [“Overview: vTM Cluster Backups” on page 237](#).

- Click **Apply**.

The required backup schedule is added to the cluster.

FIGURE 220 Cluster With Backup Schedule Added

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-RNPP-UIP9-RUA7-Q2JU	Discovered	✓	N/A	N/A		Backup Now		
Violet-Cluster	User Created		N/A	N/A		Backup Now		
Cerise-Cluster	Discovered	✓	N/A	N/A		Backup Now		
Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-05 11:00:00	Backup Now		✕

Viewing Backups for a Cluster

Once you have added a backup schedule to a cluster (see [“Adding a Backup Schedule to a Cluster” on page 242](#)), backups will begin to accumulate.

Backups are listed in the detailed view of the cluster. For example:

FIGURE 221 Cluster With Backup Schedule and Backups

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-hourly-01	2016-07-03 08:30:00	Backup Now		

Cluster Name:	<input type="text"/>
Owner:	JK ▼
Analytics Profile:	▼
Backup Schedule:	sched-hourly-01 ▼
Number of Backups:	5

Apply Revert

Backup Name	Description	Date	Retain	Actions
▶ Backup-QB4D-QTRH-N6VE-LTLZ	Cluster-AQJE-R4HV-QYRI-9F4O#74	2016-07-03 06:30		Upload Restore Compare
▶ Backup-NXIJ-BPOY-U4F6-OS28	Cluster-AQJE-R4HV-QYRI-9F4O#75	2016-07-03 07:30		Upload Restore Compare
▶ Backup-MPFV-K23G-H3FT-OVWE	Cluster-AQJE-R4HV-QYRI-9F4O#72	2016-07-03 04:30		Upload Restore Compare
▶ Backup-YFAP-9YU9-T75R-3ZAV	Cluster-AQJE-R4HV-QYRI-9F4O#71	2016-07-03 03:30	✓	Upload Restore Compare
▶ Backup-ZUOZ-TTF3-NKEZ-FJDR	Cluster-AQJE-R4HV-QYRI-9F4O#73	2016-07-03 05:30	✓	Upload Restore Compare

In this cluster:

- The cluster **Type** is *Discovered*. See [“Understanding Virtual Traffic Manager Cluster Details” on page 230](#).
- The cluster is **In Use**. That is, the cluster contains one or more vTMs.
Note: When a cluster is not **In Use**, you can delete it, see [“Deleting an Empty Virtual Traffic Manager Cluster” on page 258](#).
- The cluster does not have an assigned **Analytics Profile**. That is, analytics is not enabled on the vTMs in the cluster. See [“Configuring vTM Analytics on the Services Director” on page 281](#).
- There is a **Backup Schedule** in use on this cluster: *sched-hourly-01*
- The **Next Backup Time** for the cluster is displayed.
- The **Backup** button in the **Action** column enables you to take an immediate backup without disrupting the schedule. See [“Performing an Immediate Backup for a Cluster” on page 246](#).
- There is a listed **Owner** for the cluster.
- The maximum **Number of Backups** is 5.
- The cluster contains the three most recent backups, plus two backups that have been *retained* for future use. The retained backup will not be replaced by the addition of newer cluster backups. See [“Overview: vTM Cluster Backups” on page 237](#).

For each listed backup file:

- The default **Description** for a cluster backup is the cluster name plus a sequence number. You can update this if required, along with other details, see [“Updating Details for a Cluster Backup” on page 245](#).

- You can compare any backup to any other backup using the **Compare** button in the **Actions** column. See [“Comparing Two Cluster Backups” on page 248](#).
- You can restore any of the backups to this (or another) cluster using the **Restore** button in the **Actions** column. See [“Restoring a Backup to a Cluster” on page 250](#).
- You can upload any of the backups to any vTM using the **Upload** button in the **Actions** column. The destination vTM can be either inside or outside the cluster. You can then compare the cluster backup to either a running cluster configuration, or to another cluster backup on that vTM. See [“Uploading a Cluster Backup to a Virtual Traffic Manager” on page 253](#).

To update details for a cluster backup, see [“Updating Details for a Cluster Backup” on page 245](#).

Updating Details for a Cluster Backup

Each cluster that has an assigned backup schedule will accumulate backups over time. These backups are displayed in the detailed view of a cluster on the **vTM Clusters** page.

You cannot change the **Backup Name**, but you can update the **Description** to provide memorable information. This is useful when you choose to **Retain** a backup. See [“Overview: vTM Cluster Backups” on page 237](#).

You update details for a cluster backup from the **vTM Clusters** page.

- Access your Services Director VA from a browser, using its Service Endpoint IP Address.
- Log in as the administration user. The **Home** page appears.
- Click the **Services** menu, and then click **Services Director: vTM Clusters**. The **vTM Clusters** page appears.
- Expand the required cluster. For example:

FIGURE 222 Viewing Backups for a Cluster

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-05 11:00:00	Backup Now	Backup Now	Complete

Cluster Name:	<input type="text"/>
Owner:	JK ▼
Analytics Profile:	▼
Backup Schedule:	sched-daily-01 ▼
Number of Backups:	5
<input type="button" value="Apply"/> <input type="button" value="Revert"/>	

Backup Name	Description	Date	Retain	Actions
▶ Backup-L2CZ-IE2R-FHSD-X38R	Cluster-AQJE-R4HV-QYR1-9F4O#80	2016-07-04 10:10		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-2J4Z-RYSD-4G1R-O3T1	Cluster-AQJE-R4HV-QYR1-9F4O#81	2016-07-04 10:40		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-SU6M-PDBL-NO6Y-6L1O	Cluster-AQJE-R4HV-QYR1-9F4O#82	2016-07-04 11:00		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-FLO6-IZ01-N5TQ-D2XO	Cluster-AQJE-R4HV-QYR1-9F4O#83	2016-07-04 15:29		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-3VKN-XXUO-O33N-BKE7	Cluster-AQJE-R4HV-QYR1-9F4O#84	2016-07-04 15:30		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-QB4D-QTRH-N6VE-LTLZ	Sunday 2016/07/03	2016-07-03 06:30	✓	<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>

- Expand the required backup. For example:

FIGURE 223 Updating Details for a Cluster Backup

Backup Name	Description	Date	Retain	Actions
Backup-L2CZ-1E2R-FHSD-X38R	Cluster-AQJE-R4HV-QYR1-9F4O#80	2016-07-04 10:10	<input type="checkbox"/>	Upload Restore Compare

Description:

Retain: ☐

Apply Revert

- Update the details for the backup as required:
 - (Optional) Enter a new **Description**.
 - (Optional) Select the **Retain** check box.

Note: When a backup is *retained*, it is not deleted as newer backups are created, and does not count towards the number of backups stored by the cluster. Refer to the **Number of Backups** in step 4 and also **"Overview: vTM Cluster Backups" on page 237**.

For example:

FIGURE 224 Example: Updating Details for a Cluster Backup

Backup Name	Description	Date	Retain	Actions
Backup-L2CZ-1E2R-FHSD-X38R	Cluster-AQJE-R4HV-QYR1-9F4O#80	2016-07-04 10:10	<input checked="" type="checkbox"/>	Upload Restore Compare

Description:

Retain: ☒

Apply Revert

- Click **Apply**.

The table of backups updates to reflect the changes.

FIGURE 225 Example: Updated Cluster Backup

Backup Name	Description	Date	Retain	Actions
Backup-QB4D-QTRH-N6VE-LTLZ	Sunday 2016/07/03	2016-07-03 06:30	<input checked="" type="checkbox"/>	Upload Restore Compare
Backup-L2CZ-1E2R-FHSD-X38R	Monday 2016/07/04	2016-07-04 10:10	<input checked="" type="checkbox"/>	Upload Restore Compare ✕

Performing an Immediate Backup for a Cluster

When a cluster has an assigned backup schedule, over time it accumulates backups automatically.

However, you can also create a cluster backup at any time as an immediate manual operation.

Performing an Immediate Backup

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **Services Director: vTM Clusters**.

The **vTM Clusters** page appears.

FIGURE 226 Immediate Backup: Table of Clusters

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Backup Now		
Cerise-Cluster	Discovered	✓	N/A	N/A		Backup Now		
Cluster-RNPP-UIP9-RUA7-Q2JU	Discovered	✓	N/A	N/A		Backup Now		

4. Locate the required cluster and click the **Backup Now** button for its entry.

FIGURE 227 Immediate Backup: Starting the Operation

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Backup Now		

The Services Director attempts an immediate backup, and indicates this.

If the immediate backup succeeds, the **Last Action** and **Last Action Status** columns are updated:

FIGURE 228 Immediate Backup: Success

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Backup Now	Backup Now	Complete

If the immediate backup fails, the **Action**, **Last Action** and **Last Action Status** columns are updated:

FIGURE 229 Immediate Backup: Failure

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Clear Failed Action Retry	Backup Now	Failed

To re-attempt a failed immediate backup, see [“Retrying An Immediate Backup After a Failure” on page 247](#).

Retrying An Immediate Backup After a Failure

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **Services Director: vTM Clusters**. The **vTM Clusters** page appears.

4. Locate the required cluster. Any cluster with an immediate backup failure will show the **Action**, **Last Action** and **Last Action Status** columns as follows:

FIGURE 230 Immediate Backup: Failed Immediate Backup

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Clear Failed Action Retry Backup Now	Failed	Failed

5. Pause the pointer over the Failure warning triangle to view more information about the failure. For example:

FIGURE 231 Immediate Backup: Failed Immediate Backup Reason

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Clear Failed Action Retry Backup Now	Failed	Failed

Could not select a vTM for backup for cluster Cluster-AQJE-R4HV-QYRI-9F4O : (Error creating backup : 'Unable to access REST API Instance-OUKO-HJOV-QTGS-RRPI for creating backup.)

6. Investigate and resolve the issue. This may require you to log in to one of the vTMs in the cluster. Refer to the Virtual Traffic Manager documentation for details of vTM operations.
7. (Optional) Click the **Clear Failed Action** button for the cluster.

This action clears the **Last Action** and **Last Action Status** columns before you re-attempt the immediate backup. It is not required.

8. Once the issue is resolved, click the **Retry** button for the cluster:

FIGURE 232 Immediate Backup: Retrying an Immediate Backup

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Clear Failed Action Retry Backup Now	Failed	Failed

If the immediate backup succeeds, the failure is cleared, and the status becomes *Complete*:

FIGURE 233 Immediate Backup: Success

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Backup Now	Backup Now	Complete

If the immediate backup fails again, repeat this procedure from step 5.

Comparing Two Cluster Backups

When a cluster has an assigned backup schedule, over time it accumulates backups. Before choosing a cluster backup from which to perform a restore, it may be useful to compare two backups from the same cluster.

The resulting differences are grouped by resource type and individual resource differences.

Analyzing the differences between cluster backups supports you making an informed decision about which backup is required for a given situation.

Note: You are also able to upload a cluster backup file to a vTM, so that you can compare it to either a running cluster configuration, or to another backup on that vTM. See **“Uploading a Cluster Backup to a Virtual Traffic Manager” on page 253.**

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **Services Director: vTM Clusters**. The **vTM Clusters** page appears.
4. Expand the required cluster. The backups taken for the cluster are listed. For example:

FIGURE 234 Comparing Two Cluster Backups

Cluster Name	Type	In Use	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	sched-daily-01	2016-07-06 11:00:00	Backup Now	Backup Now	Complete

Cluster Name:
Owner:
Analytics Profile:
Backup Schedule:
Number of Backups:

Backup Name	Description	Date	Retain	Actions
▶ Backup-TYKU-FFXD-PB5M-4CG4	Cluster-AQJE-R4HV-QYR1-9F4O#87	2016-07-05 18:19		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-JROB-UMXU-C3VO-59T2	Cluster-AQJE-R4HV-QYR1-9F4O#86	2016-07-05 16:48		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-HD4G-6MJZ-ITSW-V88C	Cluster-AQJE-R4HV-QYR1-9F4O#85	2016-07-05 11:00		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-3VKN-XXUO-033N-BKE7	Cluster-AQJE-R4HV-QYR1-9F4O#84	2016-07-04 15:30		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-FLO6-IZ01-N5TQ-D2XO	Cluster-AQJE-R4HV-QYR1-9F4O#83	2016-07-04 15:29		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-L2CZ-1E2R-FHSD-X38R	Monday 2016/07/04	2016-07-04 10:10	✓	<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
▶ Backup-OB4D-QTRH-N6VE-LTLZ	Sunday 2016/07/03	2016-07-03 06:30	✓	<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>

5. Identify the first backup for the comparison and click its **Compare** button.

FIGURE 235 Identifying the First Backup

Backup Name	Description	Date	Retain	Actions
▶ Backup-TYKU-FFXD-PB5M-4CG4	Cluster-AQJE-R4HV-QYR1-9F4O#87	2016-07-05 18:19		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>

The **Compare Backups (<cluster_id>)** dialog box appears. For example:

FIGURE 236 Selecting the Second Backup

Compare Backups (Cluster-AQJE-R4HV-QYR1-9F4O)

Backup Name: Cluster-AQJE-R4HV-QYR1-9F4O#87
Compare Against:

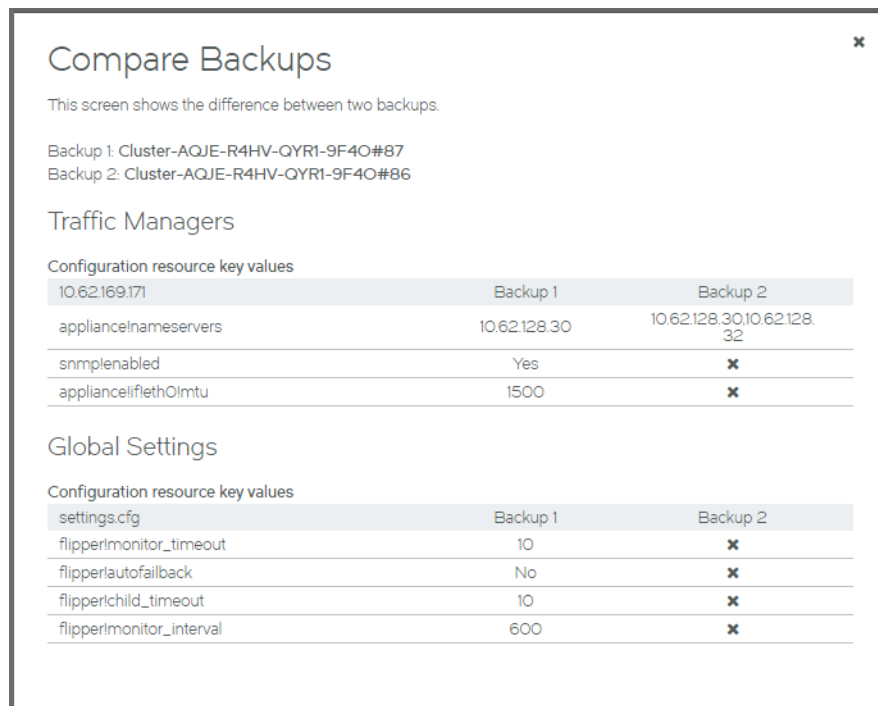
6. Select the required **Compare Against** values to identify the second backup:

- The top **Compare Against** field lists all clusters known to the Services Director. Select the current cluster (the default) or a different cluster.
- The bottom **Compare Against** field lists all backups within the selected cluster. Select the required backup for the comparison.

7. Click **Compare** to perform a comparison of the two backups.

The **Compare Backups** dialog box displays the results of the comparison. For example:

FIGURE 237 Results of the Backup Comparison



The screenshot shows a window titled "Compare Backups" with a close button (X) in the top right corner. Below the title, it says "This screen shows the difference between two backups." and lists two backups: "Backup 1: Cluster-AQJE-R4HV-QYR1-9F4O#87" and "Backup 2: Cluster-AQJE-R4HV-QYR1-9F4O#86".

Traffic Managers

Configuration resource key values

	Backup 1	Backup 2
10.62.169.171		
applianceNameservers	10.62.128.30	10.62.128.30, 10.62.128.32
snmpEnabled	Yes	✗
applianceFleethOmtu	1500	✗

Global Settings

Configuration resource key values

	Backup 1	Backup 2
settings.cfg		
flipperMonitor_timeout	10	✗
flipperAutoFailback	No	✗
flipperChild_timeout	10	✗
flipperMonitor_interval	600	✗

Backup 1 and **Backup 2** identify settings that have changed between the two backups.

Refer to the Virtual Traffic Manager documentation for details of these settings.

Restoring a Backup to a Cluster

At any point, you can restore the configuration of a cluster from a cluster backup.

Typically, the backup will be one that was generated for the cluster. However, it is possible to restore a backup from any cluster to any other cluster.

Restoring a Cluster Backup

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.

- Click the **Services** menu, and then click **Services Director: vTM Clusters**. The **vTM Clusters** page appears.
- Expand the required cluster to view its accumulated backups. For example:

FIGURE 238 Viewing Accumulated Backups for a Cluster

Cluster Name	Type	In Use	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	sched-daily-01	2016-07-06 11:00:00	Backup Now		

Cluster Name:	<input type="text"/>
Owner:	JK ▼
Analytics Profile:	▼
Backup Schedule:	sched-daily-01 ▼
Number of Backups:	5
<input type="button" value="Apply"/>	<input type="button" value="Revert"/>

Backup Name	Description	Date	Retain	Actions
Backup-QB4D-QTRH-N6VE-LTLZ	Sunday 2016/07/03	2016-07-03 06:30	✓	<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
Backup-L2CZ-1E2R-FHSD-X38R	Monday 2016/07/04	2016-07-04 10:10	✓	<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
Backup-1YKU-FFXD-PB5M-4CG4	Cluster-AQJE-R4HV-QYR1-9F4O#87	2016-07-05 18:19		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
Backup-JRO8-UMXU-C3VO-S9T2	Cluster-AQJE-R4HV-QYR1-9F4O#86	2016-07-05 16:48		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
Backup-HD4G-6MJZ-ITSW-V88C	Cluster-AQJE-R4HV-QYR1-9F4O#85	2016-07-05 11:00		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
Backup-3VKN-XXUO-033N-BKE7	Cluster-AQJE-R4HV-QYR1-9F4O#84	2016-07-04 15:30		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>
Backup-FLO6-IZ01-N5TQ-D2XO	Cluster-AQJE-R4HV-QYR1-9F4O#83	2016-07-04 15:29		<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/>

- Locate the required backup. This can be any of the listed cluster backups: scheduled, immediate or retained.

Note: If you are unsure which is required, you can compare any two backups to identify the differences, see [“Comparing Two Cluster Backups” on page 248](#).

- Click the **Restore** button for the required backup.

FIGURE 239 Starting a Restore from a Cluster Backup

Backup Name	Description	Date	Retain	Actions
Backup-QB4D-QTRH-N6VE-LTLZ	Sunday 2016/07/03	2016-07-03 06:30	✓	<input type="button" value="Upload"/> <input type="button" value="Restore"/> <input type="button" value="Compare"/> <input type="button" value="X"/>

The **Restore Backup** dialog box appears.

FIGURE 240 Selecting a Target Cluster for a Restore

Restore Backup

Backup Name: Sunday 2016/07/03

Target Cluster: Cluster-AQJE-R4H▼

- Select the **Target Cluster** from the list of clusters known to the Services Director.

8. Click **Restore**.

The Services Director begins the restore process.

FIGURE 241 Restoring a Cluster Backup: In Progress

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Backup Now	Restore	In Progress

When this completes, the selected backup has been restored to the selected cluster.

FIGURE 242 Restoring a Cluster Backup: Complete

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Backup Now	Restore	Complete

If the restore fails, the following is displayed:

FIGURE 243 Restoring a Cluster Backup: Failure

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Clear Failed Action	Restore	Failed

To resolve a failed restore, see [“Retrying A Cluster Restore After a Failure” on page 252](#).

Retrying A Cluster Restore After a Failure

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **Services Director: vTM Clusters**. The **vTM Clusters** page appears.
4. Locate the required cluster. Any cluster with a restore failure will show the **Action**, **Last Action** and **Last Action Status** columns as follows:

FIGURE 244 Cluster Restore: Failure

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Clear Failed Action	Restore	Failed

5. Pause the pointer over the Failure warning triangle to view more information about the failure. For example:

FIGURE 245 Cluster Restore: Reasons For Failure

						<div> <p>Could not find any vTMs in the cluster Cluster-AQJE-R4HV-QYR1-9F4O to restore. Uploading backup failed: Unable to access REST API Instance-OJKD-HJOV-OTGS-RRPI for uploading backup.</p> </div>		
Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Clear Failed Action	Restore	Failed

6. Investigate and resolve the issue. This may require you to log in to one of the vTMs in the cluster. Refer to the Virtual Traffic Manager documentation for details of vTM operations.
7. (Optional) Click the **Clear Failed Action** button for the cluster.

This action clears the **Last Action** and **Last Action Status** columns before you re-attempt the cluster restore. It is not required.

- Once the issue is resolved, click the **Retry** button for the cluster:

FIGURE 246 Cluster Restore: Retrying

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Clear Failed Action Retry	Restore	Failed

If the restore succeeds, the failure is cleared, and the status becomes *Complete*:

FIGURE 247 Cluster Restore: Success

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-06 11:00:00	Backup Now	Restore	Complete

If the restore fails again, repeat this procedure from step 5.

Uploading a Cluster Backup to a Virtual Traffic Manager

In addition to cluster backup comparisons (see [“Comparing Two Cluster Backups” on page 248](#)), you can upload a cluster backup file to a vTM. The uploaded cluster backup file is stored by the vTM, but not restored. This enables you to perform a comparison of the cluster backup with a running cluster configuration, or to another backup on the vTM.

After you have uploaded a cluster backup file, it is visible in the vTM's graphical user interface:

FIGURE 248 Viewing an Uploaded Cluster Backup in the vTM User Interface

BROCADE Virtual Traffic Manager Appliance Services Director 11.0

10.62.169.171 (admin/admin) Logout

Cluster: OK 0 b/s

Home Services Catalogs Diagnose Activity **System** Wizards

System: Traffic Managers Fault Tolerance Application Firewall Networking Data Plane Acceleration Sysctl Alerting

SNMP Security Users **Backups** Licenses Time Global Settings

Backups

Backup Management

View, create and import configuration backups.

Backups stored on Traffic Manager '10.62.169.171'

This section contains a list of your saved backups on this machine. To view a detailed summary, restore or export a backup click on its name.

Backup	Timestamp	Description	Compare
Backup-QB4D-QTRH-N6VE-111216 07:29	2016-07-29 07:29	SD backup	<input type="checkbox"/>
		Current Configuration	<input type="checkbox"/>

Compare

Refer to the Virtual Traffic Manager documentation for a description of supported activities with this backup.

Uploading a Cluster Backup

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **Services Director: vTM Clusters**. The **vTM Clusters** page appears.
4. Expand the required cluster to view its accumulated backups. For example:

FIGURE 249 Viewing Accumulated Backups for a Cluster

Cluster Name	Type	In Use	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	sched-daily-01	2016-07-07 11:00:00	Backup Now		

Cluster Name:	<input type="text"/>
Owner:	JK
Analytics Profile:	
Backup Schedule:	sched-daily-01
Number of Backups:	5

Apply Revert

Backup Name	Description	Date	Retain	Actions
▶ Backup-OB4D-QTRH-N6VE-LTLZ	Sunday 2016/07/03	2016-07-03 06:30	✓	Upload Restore Compare
▶ Backup-L2CZ-IE2R-FHSD-X38R	Monday 2016/07/04	2016-07-04 10:10	✓	Upload Restore Compare
▶ Backup-1YKU-FFXD-PB5M-4CG4	Cluster-AQJE-R4HV-QYRI-9F4O#87	2016-07-05 18:19		Upload Restore Compare
▶ Backup-JRO8-UMXU-C3VO-59T2	Cluster-AQJE-R4HV-QYRI-9F4O#86	2016-07-05 16:48		Upload Restore Compare
▶ Backup-HD4G-6MJZ-ITSW-V88C	Cluster-AQJE-R4HV-QYRI-9F4O#85	2016-07-05 11:00		Upload Restore Compare
▶ Backup-3VKN-XXUO-033N-BKE7	Cluster-AQJE-R4HV-QYRI-9F4O#84	2016-07-04 15:30		Upload Restore Compare
▶ Backup-FL06-IZ0I-N5TO-D2XO	Cluster-AQJE-R4HV-QYRI-9F4O#83	2016-07-04 15:29		Upload Restore Compare

5. Locate the required backup. This can be any of the listed cluster backups: scheduled, immediate or retained.

Note: If you are unsure which is required, you can compare any two backups to identify the differences, see [“Comparing Two Cluster Backups” on page 248](#).

6. Click the **Upload** button for the required backup.

FIGURE 250 Starting a Cluster Backup Upload

Backup Name	Description	Date	Retain	Actions
▶ Backup-QB4D-QTRH-N6VE-LTLZ	Sunday 2016/07/03	2016-07-03 06:30	✓	Upload Restore Compare X

The **Upload Step 1** dialog box appears.

FIGURE 251 Selecting a Cluster for an Upload

7. Select the **Target Cluster** from the list of clusters known to the Services Director.
8. Click **Next**.

The **Upload Step 2** dialog box appears.

FIGURE 252 Selecting a vTM for an Upload

9. Select the **Target Instance** from the list of vTMs for the cluster.
10. Click **Upload** to start the upload process.

When this completes, the selected cluster backup has been uploaded to the selected vTM.

FIGURE 253 Uploading a Cluster Backup: Complete

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-07 11:00:00	Backup Now	Upload	Complete

If the upload fails, the following is displayed:

FIGURE 254 Uploading a Cluster Backup: Failure

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-07 11:00:00	Clear Failed Action Retry	Upload	Failed

To resolve a failed upload, see **“Retrying A Cluster Backup Upload After a Failure”** on page 255.

Retrying A Cluster Backup Upload After a Failure

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.

- Click the **Services** menu, and then click **Services Director: vTM Clusters**. The **vTM Clusters** page appears.
- Locate the required cluster. Any cluster with an upload failure will show the **Action**, **Last Action** and **Last Action Status** columns as follows:

FIGURE 255 Cluster Upload: Failure

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-07 11:00:00	Clear Failed Action Retry	Upload	Failed

- Pause the pointer over the Failure warning triangle to view more information about the failure. For example:

FIGURE 256 Cluster Upload: Reasons For Failure

							Could not upload to Instance-R1ZQ-NEEJ-L89P-8MUP: Uploading backup failed. Unable to access REST API Instance-R1ZQ-NEEJ-L89P-8MUP for uploading backup.	
Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-07 11:00:00	Clear Failed Action Retry	Upload	Failed

- Investigate and resolve the issue. This may require you to log in to one of the vTMs in the cluster. Refer to the Virtual Traffic Manager documentation for details of vTM operations.
- (Optional) Click the **Clear Failed Action** button for the cluster.

This action clears the **Last Action** and **Last Action Status** columns before you re-attempt the cluster upload. It is not required.

- Once the issue is resolved, click the **Retry** button for the cluster:

FIGURE 257 Cluster Upload: Retrying

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-07 11:00:00	Clear Failed Action Retry	Upload	Failed

If the upload succeeds, the failure is cleared, and the status becomes *Complete*:

FIGURE 258 Cluster Upload: Success

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-AQJE-R4HV-QYRI-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-07 11:00:00	Backup Now	Upload	Complete

If the upload fails again, repeat this procedure from step 5.

Deleting a Cluster Backup

The Services Director stores the most recent cluster backups, subject to a maximum number that you can define on a per-cluster basis. Older backups beyond this maximum are deleted automatically. You can choose to mark one or more cluster backups as *retained*. *Retained* backups are not deleted automatically, and do not count towards the maximum number of backups for the cluster. See [“Updating a Virtual Traffic Manager Cluster” on page 234](#).

You can delete any cluster backup manually. To do this, expand a cluster on the **vTM Clusters** page, and locate the required cluster backup. Then, click its delete (X) button:

FIGURE 259 Deleting a Cluster Backup

Backup Name	Description	Date	Retain	Actions
▶ Backup-HD4G-6MJZ-ITSW-V88C	Cluster-AQJE-R4HV-QYR1-9F4O#85	2016-07-05 11:00		<div> <div>Upload</div> <div>Restore</div> <div>Compare</div> </div> <div>✕</div>

If you attempt to delete a *retained* cluster backup, you must confirm the deletion.

Moving a vTM Between Clusters

You cannot change a vTM's cluster from the Services Director. This is true for both registered vTMs (in *Discovered* clusters) and deployed vTMs (in *User Created* clusters).

However, you can change a vTM's cluster from the user interface of the vTM software itself. Refer to the Virtual Traffic Manager docs for information.

After you move a vTM between clusters, the existing administration credentials for the vTM in the Services Director VA will be wrong. As a result, the **Instance Health** for the vTM will change to *N/A*, and its software version will show as *Unknown*.

To fix this:

1. Access the detailed view for the vTM in the **vTM Instances** page.
2. Update the administration credentials for the vTM to those of the new cluster.

After a short time, the **Instance Health** will change to reflect the state of its new cluster, and the displayed software version will return to its usual setting.

Deleting an Empty Virtual Traffic Manager Cluster

The **vTM Clusters** page displays all clusters known to the Services Director. This page can include clusters that are not flagged as **In Use**, such as one that remains after a vTM joins another cluster, leaving its original cluster empty.

You can delete any cluster that is not flagged as **In Use**, and which does not contain cluster backups.

To delete a cluster, pause the pointer over it in the table of clusters, and then click the delete (X) button that appears at the end of the row.

FIGURE 260 Deleting an Empty Cluster

vTM Clusters

➕ Add

Cluster Name	Type	In Use	Analytics Profile	Backup Schedule	Next Backup Time	Action	Last Action	Last Action Status
Cluster-6UOX-ISLW-YBA4-II6K	Discovered		N/A	N/A		Backup Now X		
Cluster-RNPP-UIP9-RUA7-Q2JU	Discovered	✓	N/A	N/A		Backup Now		
Violet-Cluster	User Created		N/A	N/A		Backup Now		
Cerise-Cluster	Discovered	✓	N/A	N/A		Backup Now		
Cluster-AQJE-R4HV-QYR1-9F4O	Discovered	✓	N/A	sched-daily-01	2016-07-05 11:00:00	Backup Now	Backup Now	Complete

Select the **Delete** option to remove the empty cluster from the table.

Note: A dialog box appears if the empty cluster had ever contained a vTM that is now *Deleted*. This indicates that any *Deleted* vTMs will be purged from the database. For example:

FIGURE 261 Purging Deleted vTMs

Confirm Deletion

This cluster has associated instances in a 'Deleted' state, which must be purged before this cluster can be deleted. Would you like to remove them?

Instance ID	Instance Name
Instance-XMOC-A9M4-5WS5-FG2W	sienna-02
Instance-PJWS-2DIZ-RVHK-6YOK	sienna-01
Instance-2VRL-IBA3-PBSE-85QM	sienna-03

OK
Cancel

Click **OK** to purge the *Deleted* vTMs and remove the cluster from the database.

Working with User Authentication

• Overview: vTM User Authentication	259
• Overview: Services Director User Authentication	260
• Creating an Authenticator	261
• Creating a Permission Group	268
• Creating an Access Profile (vTM User Authentication Only)	272
• Applying User Authentication to a vTM	275
• Working with vTM Templates	277

The Services Director VA supports user authentication in two forms:

- vTM user authentication controls access to individual vTM instances. See **“Overview: vTM User Authentication” on page 259**.
- Services Director user authentication controls access to the Services Director's graphical user interface (GUI), command line interface (CLI) and REST API. See **“Overview: Services Director User Authentication” on page 260**.

Overview: vTM User Authentication

Each Virtual Traffic Manager (vTM) supports *user authentication*. This enables the vTM to verify the identify of any connecting user.

Note: The use of vTM user authentication is optional.

The vTM verifies a user's credentials (username and password) against two possible user authentication sources:

- Local users - user credentials are authenticated against all locally-defined user accounts (such as admin).
- Remote authenticators - user credentials are authenticated against externally-located servers that are based on RADIUS, LDAP or TACACS+ services.

Successful authentication identifies the user's permission group. This defines the activities that the connected user can perform on the vTM.

The Services Director VA enables you to optionally configure the authenticators and permissions groups that will be used by the vTMs within its estate. Specific combinations of authenticators and permission groups are combined as access profiles on the Services Director.

To configure vTM user authentication, you must create:

- One or more Services Director authenticators, see [“Creating an Authenticator” on page 261](#).
- One or more permission groups. See [“Creating a Permission Group” on page 268](#).
- One or more access profiles. See [“Creating an Access Profile \(vTM User Authentication Only\)” on page 272](#).

The Services Director Administrator chooses when to apply user authentication to a vTM. This is either:

- During the acceptance of a vTM self-registration request. See [“Accepting a Pending Self-Registration Request” on page 182](#).
- During later configuration of the vTM from the Services Director VA. See [“Applying User Authentication to a vTM” on page 275](#).

Both processes require the Services Director Administrator to choose an access profile. The access profile identifies the authenticators and permission groups that are applied to the vTM to define its user authentication. These will be applied to the vTM. All cluster members are affected.

Note: The vTM Administrator can also configure user authentication directly from the vTM. The Services Director does not track any such activity, and cannot display live user authentication settings for the vTM.

Overview: Services Director User Authentication

Services Director user authentication controls access to the Services Director's graphical user interface (GUI), command line interface (CLI) and REST API.

Note: The use of Services Director user authentication is optional.

User credentials (username and password) are evaluated against two possible user authentication sources:

- Local users - user credentials are authenticated against all locally-defined user accounts (such as admin).
- Remote authenticators - user credentials are authenticated against externally-located servers that are based on RADIUS, LDAP or TACACS+ services.

Successful authentication identifies the user's permission group. This defines the activities that the connected user can perform on the vTM.

Note: For Services Director user authentication, there is typically a single permission group, which has access to all functionality.

To configure Services Director user authentication, you must create:

- One or more Services Director authenticators, see [“Creating an Authenticator” on page 261](#).
- A permission group. See [“Creating a Permission Group” on page 268](#).

Note: Access profiles (which are required for vTM user authentication) are not required for Services Director user authentication.

Once you have created a Services Director authenticator and a permission group, the configuration of Services Director user authentication is complete.

Creating an Authenticator

Services Director supports user authentication at both the vTM level and the Services Director level.

One or more authenticators are required when establishing user authentication from the Services Director VA. An authenticator defines an external user authentication service. Three proprietary authentication services are supported, each of which has service-specific settings.

- LDAP, see [“Creating an LDAP Authenticator” on page 262](#).
- RADIUS, see [“Creating a RADIUS Authenticator” on page 264](#).
- TACACS+, see [“Creating a TACACS+ Authenticator” on page 266](#).

Authenticators are listed on the **Authenticators** page, see [“Viewing Authenticators” on page 261](#).

Note: A vTM administrator can also create and implement an authenticator on the vTM directly. Refer to the Virtual Traffic Manager documentation for details.

Viewing Authenticators

One or more authenticators are required when establishing user authentication from the Services Director VA.

The **Authenticators** page includes a table of vTM authenticators and a table of Services Director authenticators. Each entry in these tables shows the details that are common to all user authentication services (LDAP, Radius, TACACS+).

Name	Description
Authenticator Name	The name of the authenticator.
Type	The user authentication service for the authenticator. That is: LDAP, RADIUS or TACACS+.
Server	The IP address or hostname of the user authentication server.
Port	The port used to connect to the user authentication server.
Timeout	The timeout period (in seconds) for a connection to the user authentication server.
Fallback Group	The permissions group to which a valid user will belong if its group is not identified.
Status	(Services Director authenticators only). Indicates whether the authenticator is the active authenticator.

Expand an entry in either table to see full details for an authenticator. The displayed details will vary, depending on whether the authenticator is LDAP, RADIUS or TACACS+.

FIGURE 262 The Authenticators Page

Authenticators

vTM

+ Add

Authenticator Name	Type	Server	Port	Timeout	Fallback Group
▼ LDAP Server	LDAP	10.62.169.170	389	30	admin

Name:

Type:

Server:

Port:

Timeout:

Fallback Group:

Group Attribute:

Group Field:

Base DN:

Bind DN:

DN Method:

Filter:

Group Filter:

Search Password:

Search DN:

Apply

Revert

Services Director

+ Add

Authenticator Name	Type	Server	Port	Timeout	Fallback Group	Status
▶ TACACS+ Server	TACACS+	10.62.167.198	49	10	None	Enabled

Creating an LDAP Authenticator

This procedure supports both vTM authenticators and Services Director authenticators.

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **Authentication: Authenticators**. The **Authenticators** page appears.

- Click the plus symbol above either the vTM authenticator table or the Services Director authenticator table.

The **Create Authenticator** dialog box appears.

- Select the **LDAP** authenticator type, and click **Next**.

The **Create Authenticator: LDAP** dialog box appears.

FIGURE 263 Specifying LDAP Authenticator Details

Create Authenticator: LDAP

Name:	<input type="text"/>	Base DN:	<input type="text"/>
Server:	<input type="text"/>	Bind DN:	<input type="text"/>
Port:	<input type="text" value="389"/>	DN Method:	<input type="text" value="Search"/>
Timeout:	<input type="text" value="30"/> sec	Filter:	<input type="text"/>
Fallback Group:	<input type="text" value="None"/>	Group Filter:	<input type="text"/>
Group Attribute:	<input type="text"/>	Search Password:	<input type="text"/>
Group Field:	<input type="text"/>	Search DN:	<input type="text"/>

- Specify the following authenticator properties:
 - Name:** The name of the LDAP authenticator on the Services Director.
 - Server:** The IP address or hostname of the LDAP server.
 - Port:** The port used to connect to the LDAP server.
 - Timeout:** The timeout period (in seconds) for a connection to the LDAP server.
 - Group Attribute:** The LDAP attribute that gives a user's group. For example: "memberOf".

If multiple values are returned by the LDAP server the first valid one will be used.

This is required if **Fallback Group** is unset.

- Group Field** - the sub-field of the **Group Attribute** that gives a user's group.
For example: if **Group Attribute** is "memberOf" which delivers "CN=mygroup, OU=groups, OU=users, DC=mycompany, DC=local", set **Group Field** to "CN". The first matching field will be used.
- Fallback Group:** A permission group, for example: "admin".

If **Group Attribute** is not defined, or is not set for the user, the permission group named here will be used.

- **Base DN:** The base DN (Distinguished Name) for directory searches.
- **DN Method:** This value determines relevance/requirement of **Bind DN** and **Search DN**.

Use "Construct" when the bind DN for a user can be constructed from a known string. Refer to the **Bind DN** field.

Use "Search" when the bind DN for a user can be searched for in the directory. This is necessary if you have users under different directory paths. Refer to the **Search DN** and **Search Password** fields.

- **Bind DN:** A template to construct the bind DN from the username. This is only used when the **DN Method** is "Construct".

The string "%u" is replaced by the username. For example: "%u@mycompany.local" or "cn=%u, dn=mycompany, dn=local"

- **Filter:** A filter that uniquely identifies a user located under the Base DN.

The string "%u" will be substituted with the username. For example: "sAMAccountName=%u" (Active Directory), or "uid=%u" (Unix LDAP).

- **Group Filter:** If the user record returned by the LDAP **Filter** does not contain the required group information, you can specify an alternative group search filter here. This will typically be required if you have Unix/POSIX-style user records. If multiple records are returned the list of group names will be extracted from all of them.

The string "%u" will be replaced by the username. For example:
"(&(memberUid=%u)(objectClass=posixGroup))"

- **Search DN / Search Password** - the DN and password to use when searching the directory for a user's bind DN. These are only used when the **DN Method** is "Search". You can leave these blank if it is possible to perform the bind DN search using an anonymous bind.

7. (Optional) Set the **Enabled** check box if this is to be the active Services Director authenticator.
8. (Optional) Test the specified details for a Services Director user authentication by specifying a **Username** and **Password** and clicking **Test**.

Note: This function is not available for vTM authenticators.

9. Click **Finish**.

The LDAP authenticator is added to the Authenticator table.

Creating a RADIUS Authenticator

This procedure supports both vTM authenticators and Services Director authenticators.

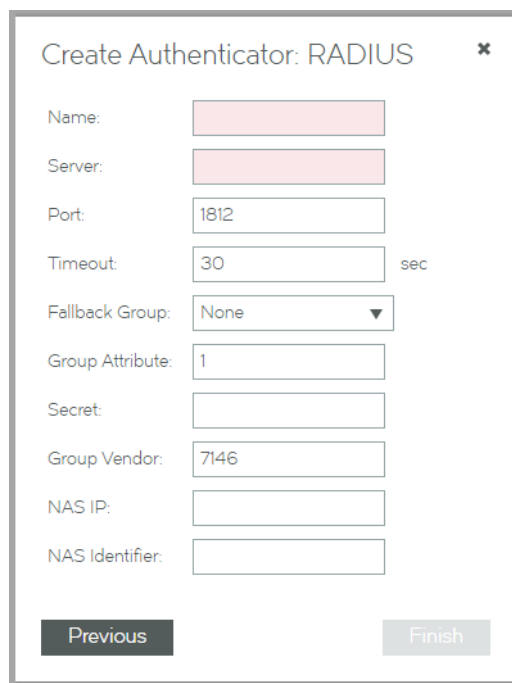
1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **Authentication: Authenticators**. The **Authenticators** page appears.
4. Click the plus symbol above either the vTM authenticator table or the Services Director authenticator table.

The **Create Authenticator** dialog box appears.

5. Select the **RADIUS** authenticator type, and click **Next**.

The **Create Authenticator: RADIUS** dialog box appears.

FIGURE 264 Specifying RADIUS Authenticator Details



Create Authenticator: RADIUS

Name:

Server:

Port:

Timeout: sec

Fallback Group:

Group Attribute:

Secret:

Group Vendor:

NAS IP:

NAS Identifier:

6. Specify the following authenticator properties:
 - **Name:** The name of the RADIUS authenticator on the Services Director.
 - **Server:** The IP address or hostname of the RADIUS server.
 - **Port:** The port used to connect to the RADIUS server.
 - **Timeout:** The timeout period (in seconds) for a connection to the RADIUS server.

- **Fallback Group:** If no group is found using the vendor and group identifiers, or the group found is not valid, the group specified here will be used.
- **Group Attribute:** The RADIUS identifier for the attribute that specifies an account's group. Optional if **Fallback Group** is specified, but required if **Fallback Group** is unset.
- **Secret:** The secret key shared with the RADIUS server.
- **Group Vendor:** The RADIUS identifier for the vendor of the RADIUS attribute that specifies an account's group.
Leave blank if using a standard attribute such as Filter-Id.
- **NAS IP:** A string identifying the Network Access Server (NAS) which is requesting authentication of the user. This value is sent to the RADIUS server.
If left blank, the address of the interface used to connect to the server will be used.
- **NAS Identifier:** The identifying IP Address of the NAS which is requesting authentication of the user. This value is sent to the RADIUS server.

7. (Optional) Set the **Enabled** check box if this is to be the active Services Director authenticator.

Note: This property is not available for vTM authenticators.

8. Click **Finish**.

The RADIUS authenticator is added to the Authenticator table.

Creating a TACACS+ Authenticator

This procedure supports both vTM authenticators and Services Director authenticators.

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **Authentication: Authenticators**. The **Authenticators** page appears.
4. Click the plus symbol above either the vTM authenticator table or the Services Director authenticator table.

The **Create Authenticator** dialog box appears.

5. Select the **TACACS+** authenticator type, and click **Next**.

The **Create Authenticator: TACACS+** dialog box appears.

FIGURE 265 Specifying TACACS+ Authenticator Details

Create Authenticator: TACACS+ ✕

Name:

Server:

Port:

Timeout: sec

Fallback Group: ▼

Group Service:

Group Field:

Secret:

Auth Type: ▼

6. Specify the following authenticator properties:

- **Name:** The name of the TACACS+ authenticator on the Services Director.
- **Server:** The IP address or hostname of the TACACS+ server.
- **Port:** The port used to connect to the TACACS+ server.
- **Timeout:** The timeout period (in seconds) for a connection to the TACACS+ server.
- **Fallback Group:** If **Group Service** is not defined, or no group value is provided for the user by the TACACS+ server, the group specified here will be used.
- **Secret:** The secret key shared with the TACACS+ server.
- **Auth Type:** The TACACS+ authentication type, either "PAP" or "ASCII".
- **Group Service:** The TACACS+ "service" that identifies a user's group field. This is required if **Fallback Group** is unset.
- **Group Field:** The TACACS+ "service" field that provides each user's group.

7. (Optional) Set the **Enabled** check box if this is to be the active Services Director authenticator.

Note: This property is not available for vTM authenticators.

8. Click **Finish**.

The TACACS+ authenticator is added to its authenticator table.

Creating a Permission Group

Services Director supports user authentication at both the vTM level and the Services Director level.

- One or more permission groups are required when establishing vTM user authentication. Each permission group defines what a user in the group can do, by combining permission names with access levels. There are four default permission groups:
 - admin - this group has full access to all vTM pages.
 - Demo - this group has full access, except to user management / system.
 - Monitoring - this group has access only to config summary / monitoring pages.
 - Guest - this group has read-only access
- A single permission group is typically required when establishing Services Director user authentication. This permission group has access to all functionality.

Permission groups are listed on the **Permission Groups** page, see [“Viewing Permission Groups” on page 268](#).

You create permission groups from the **Permission Groups** page.

- To create a permission group for vTM user authentication, see [“Creating a Permission Group \(vTM User Authentication\)” on page 270](#).
- To create a permission group for Services Director authentication, see [“Creating a Permission Group \(SD User Authentication\)” on page 272](#).

Note: The vTM administrator can create and implement a permission group on the vTM directly. Refer to the Virtual Traffic Manager documentation for details.

Viewing Permission Groups

One or more authenticators are required when establishing user authentication from the Services Director VA. Each permission group defines what a user in the group can do.

The **Permission Groups** page includes a table of permission groups for vTM user authentication, and a table of permission groups for Services Director user authentication.

FIGURE 266 The Permission Groups Page

Permission Groups

vTM

+ Add

Permission Group Name ↕	Login Timeout ↕	Description ↕
▶ admin	30	Full access to all pages
▶ Demo	30	Full access, except to user management / system
▶ Monitoring	30	Access only to config summary / monitoring pages
▶ Guest	30	Read-only access

Services Director

+ Add

Permission Group Name ↕	Description ↕
▶ admin	administration group

Each entry in the permission groups table displays summary details for the permission group.

To view full details for a vTM user authentication permission group, click the arrow on the left side of the permission group's entry.

FIGURE 267 The Permission Groups Page: vTM Permission Group

Add Permission Group

Permission Group Name:

Timeout:

Description:

Permission	None (check all)	Read Only (check all)	Full (check all)
Activity			
Connections	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Connections > Details	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Content Cache	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Content Cache > Clear	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Current Activity	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Current Activity > Edit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Download Logs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

☐ Advanced options

Add

Name	Description
Permission Group Name	The name of the permission group.
Timeout (vTM Only)	A timeout setting (in minutes) for a login session for a user in this group. A zero value indicates that sessions should never time out.
Description	A list of permissions known by the Services Director. The access level for each of these can be set to None, Read-Only or Full.
Permission	<p>A list of permissions known by the Services Director. The access level for each of these can be set to None, Read-Only or Full.</p> <p>If you click Advanced Options, you can manually specify permissions of which Services Director is not aware. That is, you can reference any permission that is supported by the vTM. To find these permission names, refer to the Virtual Traffic Manager documentation.</p> <p>Note: The Services Director VA does not verify permissions entered under Advanced Options. The vTM itself verifies all permissions when the permission group is applied to the vTM. Any permission that is not recognized by the vTM is ignored.</p>

To view full details for a Services Director user authentication permission group, click the arrow on the left side of the permission group's entry.

FIGURE 268 The Permission Groups Page: Services Director Permission Group

The screenshot shows the 'admin' permission group selected in a table. Below the table, the details for the 'admin' group are displayed: 'Permission Group Name' is 'admin', 'Description' is 'administration group', and 'Permissions' are 'Full administrator permissions'. At the bottom, there are 'Apply' and 'Revert' buttons.

Typically, there is only one Services Director user authentication permission group.

Creating a Permission Group (vTM User Authentication)

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **Authentication: Permission Groups**. The **Permission Groups** page appears.
4. Click the plus symbol above the vTM permission group table.

The **Add Permission Group** dialog box appears.

FIGURE 269 Adding a Permission Group: vTM User Authentication

Add Permission Group [X]

Permission Group Name:

Timeout:

Description:

Permission	None (check all)	Read Only (check all)	Full (check all)
Activity			
Connections	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Connections > Details	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Content Cache	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Content Cache > Clear	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Current Activity	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Current Activity > Edit	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Download Logs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

☐ Advanced options

5. Specify a **Permission Group Name**.
6. Specify a **Timeout** period, in minutes.
7. (Optional) Add a description for the permission group.
8. Specify an access level for each listed **Permission**. That is, None, Read-Only or Full.
 - To select None for all listed permissions, click **None (check all)**.
 - To select Read-Only for all listed permissions, click **Read-Only (check all)**.
 - To select Full for all listed permissions, click **Full (check all)**.
9. To specify a permission for an unlisted **Permission**:
 - a. Click **Advanced Options**.
 - b. Enter the name of the **Permission**. You can reference any permission that is supported by the vTM. To find these permission names, refer to the Virtual Traffic Manager documentation.
 - c. Select the required access level. That is, None, Read-Only or Full.
10. Click **Add** to create the vTM permission group.

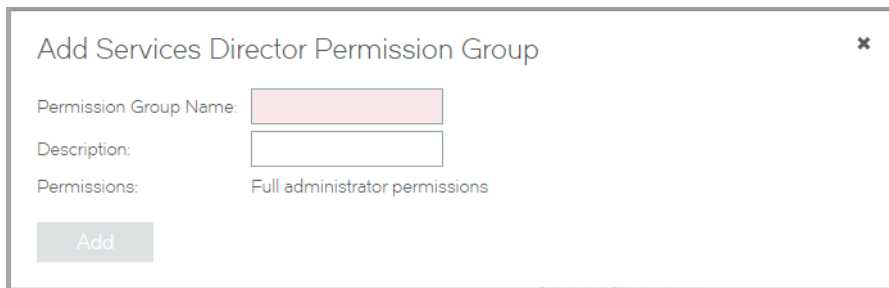
Note: The vTM administrator can create and implement a permission group on the vTM. Refer to the Virtual Traffic Manager documentation for details.

Creating a Permission Group (SD User Authentication)

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **Authentication: Permission Groups**. The **Permission Groups** page appears.
4. Click the plus symbol above the Services Director permission group table.

The **Add Permission Group** dialog box appears.

FIGURE 270 Adding a Permission Group: Services Director User Authentication



5. Specify a **Permission Group Name**.
6. (Optional) Add a description for the permission group.
7. Click **Add** to create the Services Director permission group.

Creating an Access Profile (vTM User Authentication Only)

An access profile is required when establishing user authentication for a vTM from the Services Director VA. An access profile combines an authenticator with one or more permission groups. When an access profile is selected, the authenticator and permission groups included in the profile are used by the vTM to define its user authentication.

Note: Access profiles are not required when creating Services Director user authentication.

Access profiles are listed on the **Access Profiles** page, see [“Viewing Access Profiles” on page 273](#).

You create access profiles from the **Access Profiles** page, see [“Creating an Access Profile” on page 274](#).

Note: The use of access profiles enable the Services Director Administrator to set the user authentication on the vTM from the Services Director VA. However, the vTM Administrator can also configure user authentication directly from the vTM. The Services Director does not track any such activity, and cannot display live user authentication settings for the vTM.

Viewing Access Profiles

An access profile is required when establishing user authentication for a vTM from the Services Director VA. An access profile combines an authenticator with one or more permission groups. When it is selected, the authenticator and permission groups included in the access profile are used by the vTM to define its user authentication.

Note: Access profiles are not supported for Services Director user authentication.

The **Access Profiles** page shows a table of all access profiles defined on the Services Director. Each entry in the table shows summary details for an access profile.

Name	Description
Access Profile Name	<p>The name of the access profile. This is used when applying an access profile to:</p> <ul style="list-style-type: none"> a <i>Pending</i> self-registration request by a vTM. See “Accepting a Pending Self-Registration Request” on page 182. one or more registered/deployed vTMs. See “Applying User Authentication to a vTM” on page 275.
Authenticator	The selected authenticator for the access profile. See “Creating an Authenticator” on page 261 .
Permission Groups	A list of permission groups included in the access profile. There are four default permission groups, but you can define others. See “Creating a Permission Group” on page 268 .
Actions	The Apply to vTM Instance(s) control in this column enables you to apply the permissions groups and authenticators associated with this access profile to one or more vTMs. See “Applying User Authentication to a vTM” on page 275 .

To view full details for an access profile, click the arrow on the left side of the access profile’s entry.

FIGURE 271 The Access Profiles Page

Access Profiles

+ Add

Access Profile Name	Authenticator	Permission Groups	Actions
▶ LDAP All	LDAP Server	Full access to all pages	<button>Apply</button>
▼ LDAP Statistics	LDAP Server	Read-only access	<button>Apply</button>

Name:

Authenticator:

Permission Groups:

Permissions Group	Include?
admin	<input type="checkbox"/>
Demo	<input type="checkbox"/>
Monitoring	<input type="checkbox"/>
Guest	<input checked="" type="checkbox"/>

Apply Revert

▶ RADIUS All	RADIUS Server	Full access to all pages	<button>Apply</button>
▶ TACACS+ All	TACACS+ Server	Full access to all pages	<button>Apply</button>

Creating an Access Profile

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **Authentication: Access Profiles**. The **Access Profiles** page appears.
4. Click the plus symbol above the access profile table.

The **Add Access Profile** dialog box appears.

FIGURE 272 Adding an Access Profile

Add Access Profile [X]

Access Profile Name:

Authenticator: TACACS+ Server ▼

Permission Groups:	Permissions Group	Include?
	admin	<input type="checkbox"/>
	Demo	<input type="checkbox"/>
	Monitoring	<input type="checkbox"/>
	Guest	<input type="checkbox"/>

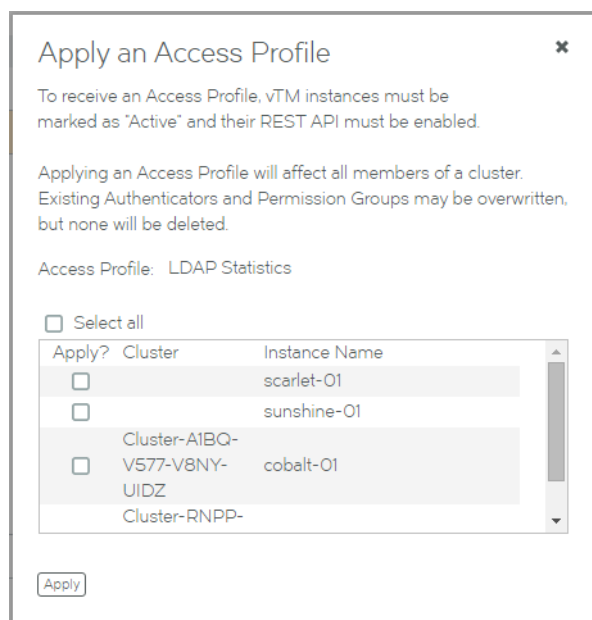
5. Specify an **Access Profile Name**.
6. Select an **Authenticator**.
7. Select one or more permission groups.
8. Click **Add** to create the access profile.

Applying User Authentication to a vTM

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Catalogs** menu, and then click **Authentication: Access Profiles**. The **Access Profiles** page appears.
4. In the table of access profiles, locate the required access profile. Expand the entry to confirm its properties if required.
5. Click the **Apply** button that is next to the required access profile.

The **Apply an Access Profile** dialog box appears. This dialog box lists all vTMs that are *Active* and with a REST API enabled.


FIGURE 273 Applying an Access Profile



6. Select the check box for each required vTM instance, or click **Select All**.
7. Click **Apply**.

A summary of selections appears. For example:


FIGURE 274 User Authentication Changes: Summary



8. Click **OK** to continue.


The permissions groups and authenticators associated with the chosen access profile are applied to the selected vTMs. A progress bar tracks this:

FIGURE 275 User Authentication Changes: Progress



Once the changes are complete, a message appears:

FIGURE 276 User Authentication Changes: Complete



9. Click **OK**. The process is complete.

Working with vTM Templates

During the process of configuring a vTM for self-registration, you can mark a vTM as a template vTM. This prevents it from self-registering, but ensures that all vTMs made from the template will request self-registration.

The template vTM is visible in the list of virtual machines in VMware, and can be used to create other vTMs. Refer to the Virtual Traffic Manager documentation.

Working with vTM Analytics

- **Overview: vTM Analytics (Enterprise Customers Only)** 279
- **Creating Analytics Resources** 285
- **Enabling Analytics on a vTM Cluster** 297
- **Working with Analytics Data on the Services Director** 298

Overview: vTM Analytics (Enterprise Customers Only)

Services Director supports the configuration and activation of analytics data export on a cluster of Virtual Traffic Managers (vTMs). Each vTM operating at version 17.2 or later supports vTM Analytics. vTM Analytics enables a vTM to send analytics data to an Analytics System.

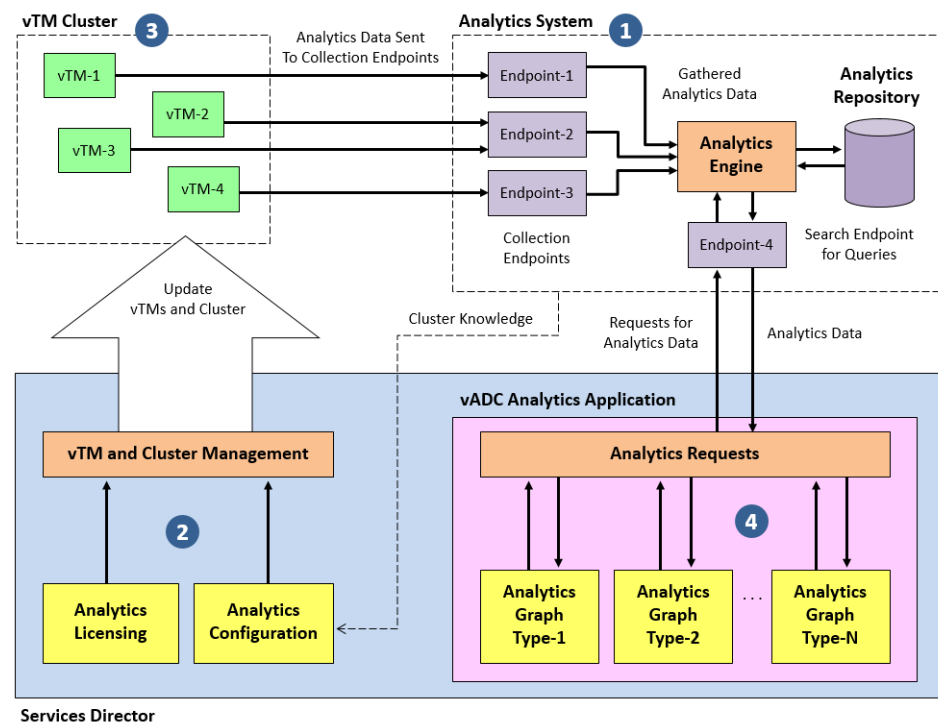
Collected data can be queried using the **vADC Analytics** application that is embedded in the graphical user interface of the Services Director VA. This displays tailored graphical reports about the vTMs in its estate.

Note: The use of vTM Analytics is optional, and is only available to customers who purchase an Analytics Resource Pack license.

Note: Currently, the **vADC Analytics** application is best supported by the Google Chrome browser.

The vTM Analytics process operates as follows:

FIGURE 277 vTM Analytics Overview



1. Outside of Services Director and the Virtual Traffic Manager, you must install and configure an Analytics System. See [“Understanding the Analytics System” on page 280](#).

Note: Services Director currently supports retrieval of analytics data from the Splunk®¹ platform only.

2. On the Services Director, you install an Analytics Resource Pack License, and create all required analytics resources. These are then used to prepare both the cluster and its vTMs for the production of analytics data. See [“Configuring vTM Analytics on the Services Director” on page 281](#).
3. The vTMs in the cluster, now configured to export analytics data, begin to transmit analytics data to the Analytics System, subject to available bandwidth in the Analytics Resource Pack license. See [“Understanding the Automatic Export of vTM Analytics Data” on page 283](#).
4. On the Services Director, the vADC Analytics Application can then query the Analytics System to present the data as a variety of analytics graphs. See [“Querying vTM Analytics from the Services Director” on page 284](#).

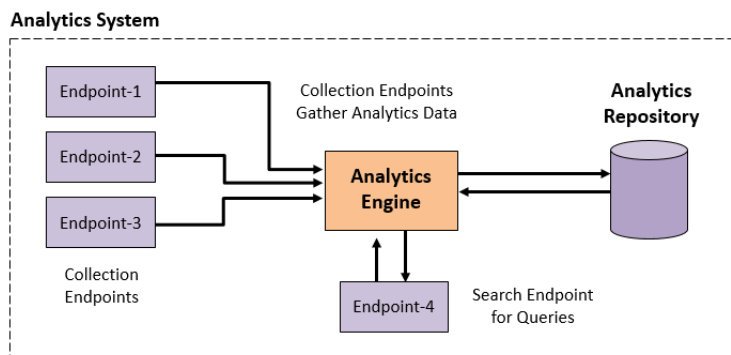
Understanding the Analytics System

The vTM Analytics functionality requires an operational *Analytics System*.

An Analytics System is a grouping of third-party machines, virtual machines, ports, repositories and software that operates collectively to collate analytics data and deliver the required analytics capability.

Note: Currently, the Services Director supports analytics using the Splunk platform.

FIGURE 278 vTM Analytics Overview: A Generalized Analytics System



This diagram is generalized; the creation, configuration and operation of the Analytics System will be tailored to your network. These activities are outside the scope of both the Services Director and the Virtual Traffic Manager products.

1. Splunk is a registered trademark of Splunk Inc. in the USA and other countries.

In general terms, your Analytics System will include:

- An analytics repository to store analytics data.
- An analytics engine that controls the collection, storage and retrieval of analytics data.
- One or more Collection Endpoints. Each collection endpoint receives analytics data from one or more vTMs, including transaction metadata and log data. Typically there will be multiple collection endpoints. Each of these endpoints must be recorded as a Collection Endpoint resource on the Services Director, see [“Adding a Collection Endpoint Resource to the Services Director” on page 289](#).
- One Search Endpoint. This unique endpoint is used by the Services Director to perform queries against analytics data stored in the analytics repository. This endpoint must be recorded as a Search Endpoint resource on the Services Director, see [“Adding a Search Endpoint Resource to the Services Director” on page 293](#).

Once the Analytics System is ready, you can use the Services Director to license and configure vTM analytics data export, see [“Configuring vTM Analytics on the Services Director” on page 281](#).

Configuring vTM Analytics on the Services Director

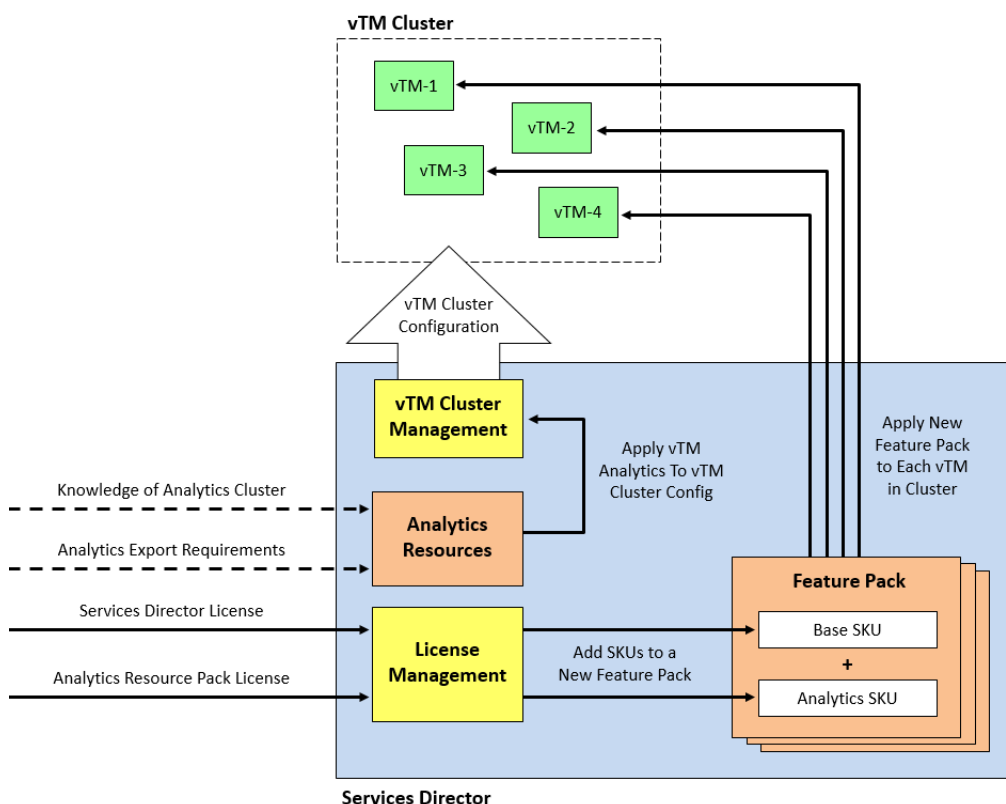
Before you can configure analytics data export on the vTMs in the estate of the Services Director, you must add an Analytics Resource Pack License to the Services Director, and create all required resources on the Services Director. To do this, you need knowledge of the Analytics System implementation. Specifically, the required endpoints and URLs.

- An Analytics Resource Pack License is required to enable analytics on a fixed number of vTMs. This license defines how many vTMs can be configured to export analytics data to the Analytics System. You must add this to the licenses on the Services Director, see [“Adding a License to the Services Director” on page 128](#).
- Feature Pack resources, each of which references both a Services Director base SKU and an ENT-ANALYTICS add-on SKU. These SKUs are enabled by the Analytics Resource Pack License above. See [“Adding a Feature Pack to the Services Director” on page 129](#).
- Log Export Type resources, each of which identifies the log types that will be exported by the vTM. See [“Creating a Log Export Type” on page 285](#).
- Analytics Profile resources, each of which identifies the types of analytics data (transaction data and logs) exported by the vTM. See [“Creating an Analytics Profile” on page 287](#).
- Collection/Search Endpoint resources, each of which identifies an endpoint in the Analytics System. A single Search Endpoint resource defines where the Services Director will direct queries to in the Analytics System, and a pool of Collection Endpoint resources defines where analytics data will be exported to by the vTMs to the Analytics System. See [“Adding Analytics Endpoint Resources to the Services Director” on page 289](#).

Once the Analytics Resource Pack License and the required resources are in place, you can configure analytics on the Services Director and the vTMs in its estate. To do this, you require:

- A single new Feature Pack for all of the vTMs in the vTM cluster. This must include both a Services Director base SKU and an ENT-ANALYTICS add-on SKU.
- An Analytics Profile to identify the analytics data that will be exported to the Analytics System by the vTMs.

FIGURE 279 vTM Analytics Overview: Configuring Services Director and its vTMs



You must then update all vTMs in the cluster to use the new Feature Pack. See [“Applying a Feature Pack to Registered Instances” on page 140](#).

You can then enable analytics on all vTMs in a cluster by applying the required analytics profile to the cluster. You do this from the **vTM Clusters** page. See [“Enabling Analytics on a vTM Cluster” on page 297](#).

Each vTM is assigned an analytics Collection Endpoint automatically by the Services Director from its pool of Endpoints.

Note: The maximum number of vTMs that can be licensed to produce analytics data is limited only by the available analytics bandwidth in the Analytics Resource Pack License. You can add additional Analytics Resource Pack Licenses to increase this maximum.

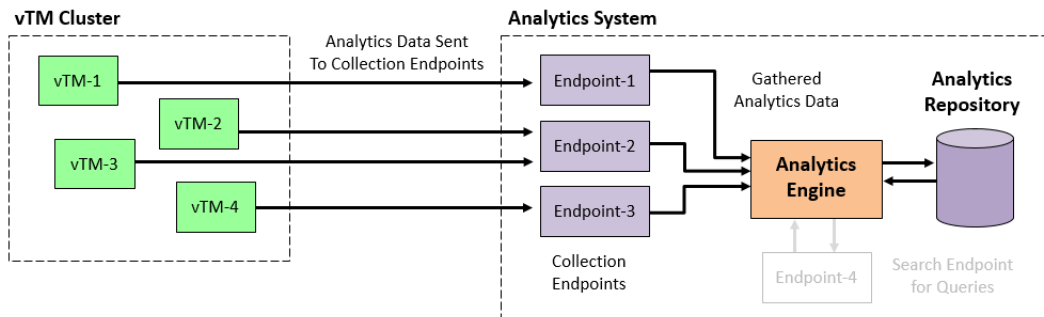
Note: Services Director applies the analytics configuration to a single vTM, and vTM cluster replication ensures it reaches all the members of the cluster.

After this process completes, all vTMs in the cluster are configured and licensed for vTM Analytics, and the export of analytics data begins. See [“Understanding the Automatic Export of vTM Analytics Data” on page 283](#).

Understanding the Automatic Export of vTM Analytics Data

Once all vTMs in the cluster are configured and licensed for vTM Analytics (see [“Configuring vTM Analytics on the Services Director” on page 281](#)), export of analytics data begins.

FIGURE 280 vTM Analytics Overview: Gathering vTM Analytics Data



Each vTM transmits the content defined by the cluster’s analytics profile to its assigned collection endpoint on the Analytics System. This data is processed and stored in the analytics repository.

Note: The transmission and processing of analytics data between the vTMs and the Analytics System is outside the scope of Services Director. Refer to the Virtual Traffic Manager documentation.

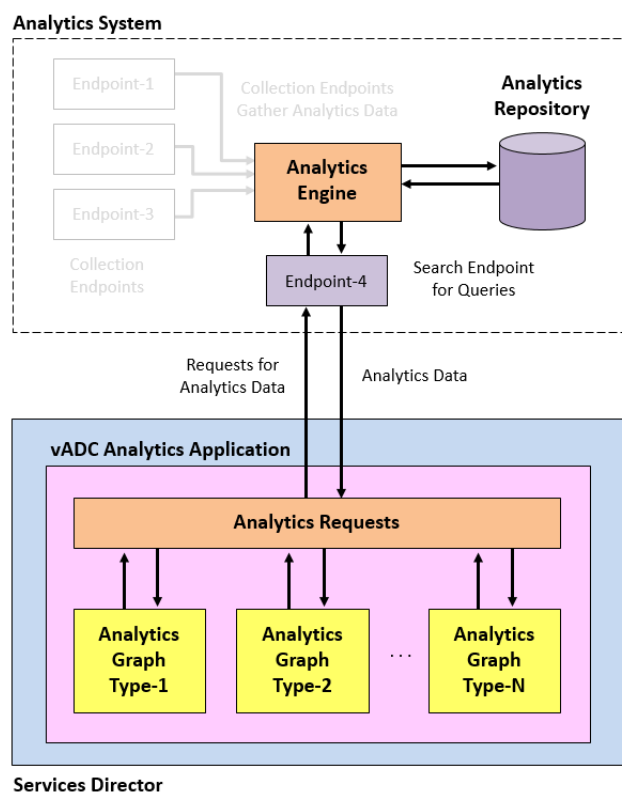
Once the Analytics Repository starts to accumulate data, the data can be queried by the embedded **vADC Analytics** application on the Services Director. See [“Querying vTM Analytics from the Services Director” on page 284](#).

Querying vTM Analytics from the Services Director

Analytics data that is stored in an Analytics System can be queried and retrieved by the embedded **vADC Analytics** application on the Services Director to enable a number of graphical analytics reports. The requests are driven from the user interface for each graph type, and sent to the Search Endpoint for the Analytics System from the Services Director. The retrieved results are displayed within the graphs on the Services Director, and can then be filtered, drilled into, and analyzed. See [“Configuring vTM Analytics on the Services Director” on page 281](#).

Note: Querying of an Analytics System can be performed by all customers who configure a Search Endpoint.

FIGURE 281 vTM Analytics Overview: Querying Analytics Data



Creating Analytics Resources

After you have added the required Analytics Resource Pack License to the Services Director, you must create the required resources on the Services Director:

- Create a new Feature Pack that includes both a base SKU and a resource SKU that supports vTM analytics. See [“Adding a Feature Pack to the Services Director” on page 129](#).
- Create one or more Log Export Type resources, each of which identifies the log types that will be exported by the vTM. See [“Creating a Log Export Type” on page 285](#).
- Create one or more Analytics Profile resources, each of which identifies the types of analytics data (transaction data and logs) exported by the vTM. See [“Creating an Analytics Profile” on page 287](#).
- Collection/Search Endpoint resources, each of which identifies an endpoint in the Analytics System:
 - A single search Endpoint is always used for Services Director queries.
 - All other Endpoints are used for data collection. All defined collection Endpoints are handled as a single pool by the Services Director, and allocated to vTMs automatically. See [“Adding Analytics Endpoint Resources to the Services Director” on page 289](#).

Creating a Log Export Type

The **Log Export Types** page lists all existing log export types in a table. Each entry identifies one or more files that will be sent to the Analytics System by the vTM.

Note: You combine Log Export Types with transaction settings to form an Analytics Profile, see [“Creating an Analytics Profile” on page 287](#).

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The Home page appears.
3. Click the **Catalogs** menu, and then click **Analytics > Log Export Types**.

The **Log Export Types** page appears. By default, a number of key Log Export Types are installed with the product. The default Log Export Types may be sufficient for your analytics requirements.

FIGURE 282 Analytics: Log Export Types

Log Export Types

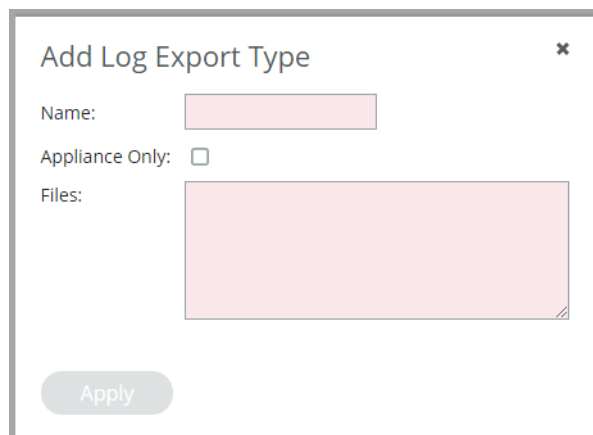
+ Add

Name	ID	Files	Appliance Only
▶ Admin Server Access	Admin Server Access	%ZEUSHOME%/admin/log/access*	
▶ Application Firewall	Application Firewall	%ZEUSHOME%/zxtm/log/stingrayafm/log-master/* %ZEUSHOME%/zxtm/log/stingrayafm/log/*	
▶ Audit Log	Audit Log	%ZEUSHOME%/zxtm/log/audit*	
▶ Data Plane Acceleration	Data Plane Acceleration	%ZEUSHOME%/zxtm/log/dpa_errors*	✓
▶ Event Log	Event Log	%ZEUSHOME%/zxtm/log/errors*	
▶ Process Monitor	Process Monitor	%ZEUSHOME%/zxtm/log/procmon*	
▶ Routing Software	Routing Software	%ZEUSHOME%/zxtm/log/routing_sw*	✓
▶ System - authentication log	System - authentication log	/var/log/auth.log*	✓
▶ System - syslog	System - syslog	/var/log/syslog*	✓

- Click the **Add** button above the **Log Export Types** table.

The **Add Log Export Type** dialog box appears.

FIGURE 283 Analytics: Adding a Log Export Type



The dialog box titled "Add Log Export Type" contains the following fields:

- Name:** A text input field.
- Appliance Only:** A checkbox.
- Files:** A large text area for specifying file paths.
- Apply:** A button at the bottom left.

- Enter a **Name** for the Log Export Type.

This name will appear in the **Log Export Types** table.

- (Optional) Select the **Appliance Only** check box if this is only supported on Virtual Appliance installations of the vTM, and not on software installations.

7. Enter one or more file names or directories as **Files**.

- Where you want to specify more than one entry, use a space-separated list.
- The asterisk wild card is supported for multiple selections. For example:

```
/var/log/auth.log*
```

- The `%ZEUSHOME%` system variable enables you to specify file structures relative to the vTM's home directory. For example:

```
%ZEUSHOME%/admin/log/access*
```

8. Click **Apply**. The new Log Export Type is added to the **Log Export Types** table.

9. Repeat this process to create all required Log Export Types.

You must then combine one or more Log Export Types with transaction settings to form an Analytics Profile. See [“Creating an Analytics Profile” on page 287](#).

Creating an Analytics Profile

The **Analytics Profiles** page lists all existing Analytics Profiles in a table. Each entry identifies the Log Export Types and transactions settings that will be sent to the Analytics System by a vTM that uses the Analytics Profile.

Note: You must create all required Log Export Types before you begin, see [“Creating a Log Export Type” on page 285](#).

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The Home page appears.
3. Click the **Catalogs** menu, and then click **Analytics > Analytics Profiles**.

The **Analytics Profiles** page appears.

FIGURE 284 Analytics: Analytics Profiles

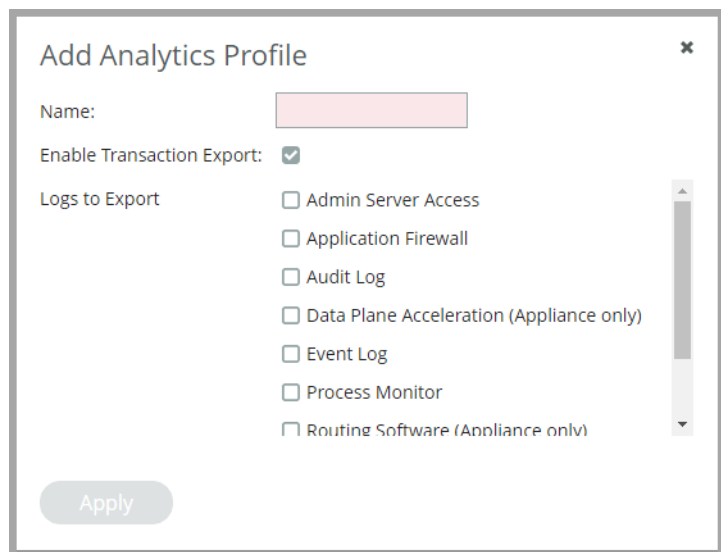
Analytics Profiles

+ Add				
Name ▾	ID ▾	Logs to export	Transaction Data Export	
▶ Audit Only	Analytics-Profile-S98X-8LJE-0Z82-7NJJ	Audit Log	Enabled	
▶ Event Only	Analytics-Profile-4WMJ-0MQB-NVEH-LX2L	Event Log	Enabled	

4. Click the **Add** button above the **Analytics Profiles** table.

The **Add Analytics Profile** dialog box appears.

FIGURE 285 Analytics: Creating an Analytics Profile



Add Analytics Profile [X]

Name:

Enable Transaction Export: ☒

Logs to Export

- ☐ Admin Server Access
- ☐ Application Firewall
- ☐ Audit Log
- ☐ Data Plane Acceleration (Appliance only)
- ☐ Event Log
- ☐ Process Monitor
- ☐ Routing Software (Appliance only)

Apply

5. Enter a **Name** for the Analytics Profile.

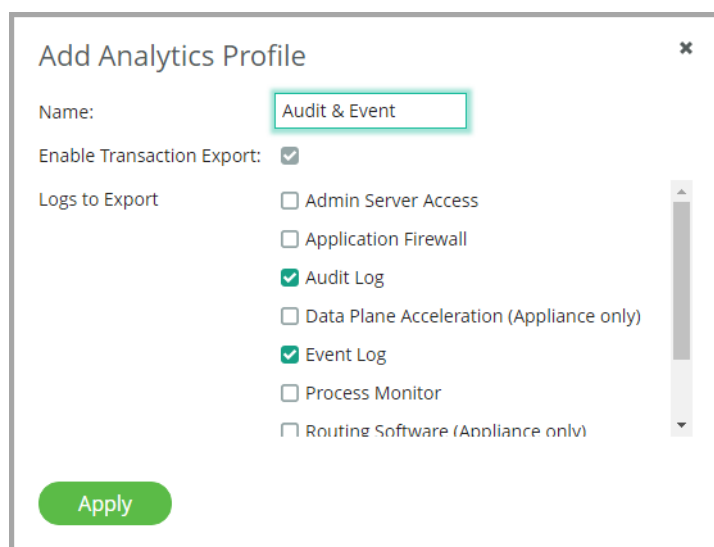
This name will appear in the **Analytics Profiles** table.

6. Select the **Enable Transaction Export** check box to include transaction metadata in the Analytics Profile.

Note: By default, transaction metadata is exported along with any selected logs. If you do not want to export transaction metadata, clear the **Enable Transaction Export** check box.

7. Select the check box for each required Log Export Type from the **Logs to Export** list. For example:

FIGURE 286 Analytics: Specifying an Analytics Profile



Add Analytics Profile [X]

Name:

Enable Transaction Export: ☒

Logs to Export

- ☐ Admin Server Access
- ☐ Application Firewall
- ☒ Audit Log
- ☐ Data Plane Acceleration (Appliance only)
- ☒ Event Log
- ☐ Process Monitor
- ☐ Routing Software (Appliance only)

Apply

Note: Where a Log Export Type is supported on Virtual Appliance installations of the vTM, this is indicated. For example, the *Data Plane Acceleration (Appliance only)* Log Export Type. When a Log Export Type is applied to a software vTM, any “Appliance only” Log Export Types are ignored.

8. Click **Apply**. The new Analytics Profile is added to the **Analytics Profiles** table.
9. Repeat this process to create all required Analytics Profiles.

Once you have created all required resources, you can apply an Analytics Profile to one or more vTM clusters. See [“Enabling Analytics on a vTM Cluster” on page 297](#).

Adding Analytics Endpoint Resources to the Services Director

Before you can configure analytics on the vTMs in the estate of the Services Director, you must create an Endpoint resource for each of the endpoints on the Analytics System. This includes:

- A pool of Collection Endpoint resources, each of which describes a collection endpoint in the Analytics System that is used to gather analytics data from the vTM cluster. See [“Adding a Collection Endpoint Resource to the Services Director” on page 289](#).
- A Search Endpoint resource. The endpoint identified by this resource is used by the Services Director to perform queries against gathered analytics data in the Analytics System. See [“Adding a Search Endpoint Resource to the Services Director” on page 293](#).

Adding a Collection Endpoint Resource to the Services Director

A collection endpoint is an element of the Analytics System. Each collection endpoint receives analytics data from one or more vTMs. See [“Understanding the Automatic Export of vTM Analytics Data” on page 283](#).

You must add a Collection Endpoint resource to the Services Director for each collection endpoint in the Analytics System. The Services Director maintains a pool of these resources, and references them when you configure analytics on a vTM cluster from the Services Director.

The **Analytics Endpoints** page lists all existing Collection Endpoint resources in a table.

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The Home page appears.
3. Click the **Catalogs** menu, and then click **Analytics > Analytics Endpoints**.

The **Analytics Endpoints** page appears.

FIGURE 287 Analytics: Analytics Endpoints Page

Analytics Endpoints

Collection Endpoints

+ Add

Name	ID	Transaction Export	Log Export
No Data			

- Click the **Add** button above the **Collection Endpoints** table.

The **Add Collection Endpoint** dialog box appears.

FIGURE 288 Analytics: Adding a Collection Endpoint

Add Collection Endpoint

Name:

Transaction Export Collector Settings

Address (<IP address/hostname>:<port>):

Export over TLS:

Verify TLS:

Certificate:

From file

Choose File

From text

Log Export Collector Settings

HTTP(S) URL:

Verify TLS:

Authentication Method:

None

Certificate:

From file

Choose File

From text

Apply

5. Enter a **Name** for the Collection Endpoint resource.

This name will appear in the **Collection Endpoints** table.

6. If the collection endpoint will accept transaction metadata, you must now define the **Transaction Export Collector Settings** for its resource:

- Enter an **Address** for the collection endpoint in the Analytics System. This takes the form:
`<IP address/hostname>:<port>`
 You cannot specify a protocol or a filepath.
- If you want Transport Layer Security (TLS) to be used during transaction metadata export, select the **Export over TLS** check box.
- If the **Export over TLS** check box is selected, you can choose to verify the TLS connection by selecting the **Verify TLS** check box.
- If the **Export over TLS** check box is selected, you must provide an SSL **Certificate**. To do this, either browse for the required certificate file in the **From file** property, or paste the contents of the certificate into the **From text** property.

7. If the Collection Endpoint will accept log data, you must now define the **Log Export Collector Settings** for its resource:

- Enter an **HTTP(S) URL** for the collection endpoint in the Analytics System. This takes the form:

`<protocol><server>:<port><filepath>`
 The protocol can be either *http://* or *https://*.

Note: If you want Transport Layer Security (TLS) to be used during data export, use the *https://* protocol.

- If TLS is used, you can choose to verify the TLS connection by selecting the **Verify TLS** check box.
- If TLS is used, you must provide an SSL **Certificate**. To do this, either browse for the required certificate file in the **From file** property, or paste the contents of the certificate into the **From text** property.
- Select the required **Authentication Method**:
 - "None". If you select this option, no additional authentication properties are required.
 - "Basic HTTP Authentication". If you select this option, you must then specify a **Username** and **Password**.
 - "Splunk". If you select this option, you must then specify the **HEC Token** from the Splunk platform.

FIGURE 289 Analytics: Collection Endpoint Example

Add Collection Endpoint

Name: JK-EP-Collection-01

Transaction Export Collector Settings

Address (<IP address/hostname>:<port>): demo.com:7070

Export over TLS: ☒

Verify TLS: ☒

Certificate:

☐ From file [Choose File](#)

☒ From text

```
-----BEGIN PRIVATE KEY-----
MIICdwIBADANBgkqhkiG9w0BAQEFAASCAmEwg
LslHOqhF4XoX7au5Fe4B52h7Jam1F5u8G+Q0pja
Squ1qwYyOi3a2GLicugm4it/jHkUybeVWoz5blej
9BivF+/6tMChFonT4RHjxGrfWB8vAgMBAAECgYB
O9ZIM7nILyW5seje1QUQ/WxUklqm12f+NUpl4A
-----
```

Log Export Collector Settings

HTTP(S) URL: https://demo.com:8080/logs/collector

Verify TLS: ☒

Authentication Method: Basic HTTP Authen ▼

Username: admin

Password:

Certificate:

☐ From file [Choose File](#)

☒ From text

```
-----BEGIN PRIVATE KEY-----
MIICdwIBADANBgkqhkiG9w0BAQEFAASCAmEwg
LslHOqhF4XoX7au5Fe4B52h7Jam1F5u8G+Q0pja
Squ1qwYyOi3a2GLicugm4it/jHkUybeVWoz5blej
9BivF+/6tMChFonT4RHjxGrfWB8vAgMBAAECgYB
O9ZIM7nILyW5seje1QUQ/WxUklqm12f+NUpl4A
-----
```

Apply

- Click **Apply**. The new Collection Endpoint resource is added to the **Collection Endpoints** table.

FIGURE 290 Analytics: Added Collection Endpoint

Collection Endpoints

⊕ Add

Name	ID	Transaction Export	Log Export
▶ JK-EP-Collection-01	Collection-Endpoint-S05F-7FSJ-L45E-6L51	demo.com:7070	https://demo.com:8080/logs/collector

- (Optional) Expand the Collection Endpoint resource entry to view its full details.
- Repeat this process to create all required Collection Endpoint resources.

You must also create a single Search Endpoint resource. See **"Adding a Search Endpoint Resource to the Services Director" on page 293**.

Adding a Search Endpoint Resource to the Services Director

A search endpoint is an element of the Analytics System. The search endpoint receives analytics queries from the Services Director, and returns analytics data to the Services Director. See [“Understanding the Automatic Export of vTM Analytics Data” on page 283](#).

You must add a single Search Endpoint resource to the Services Director to record the properties of the Analytics System's search endpoint.

Note: Querying of an Analytics System can be performed by any customer who configures a Search Endpoint.

Note: Multiple Search Endpoint resources are not supported.

The **Analytics Endpoints** page displays the **Search Endpoints** table.

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The Home page appears.
3. Click the **Catalogs** menu, and then click **Analytics > Analytics Endpoints**. The **Analytics Endpoints** page appears, which includes a table of Search Endpoints.

FIGURE 291 Analytics: Analytics Endpoints Page

Search Endpoints

 Add

Name	ID	Address
No Data		

- Click the **Add** button above the **Search Endpoints** table. The **Add Search Endpoint** dialog box appears.

FIGURE 292 Analytics: Adding a Search Endpoint

Add Search Endpoint [X]

Name:

Address:

Transactions index:

Logs index:

Query using TLS: ☐

Certificate: ☒ From file [Choose File](#)

☒ From text

Username:

Password:

- Enter a **Name** for the Search Endpoint resource.
This name will appear in the **Search Endpoints** table.
- Enter an Address for the search endpoint in the Analytics System. This takes the form:

`<server>:<port>`

You cannot specify a protocol or a filepath.

Note: You can test the connection to this address later in this procedure.

- Specify the **Transactions Index**. This is the index used to store transaction data on the Splunk platform. For example, *zxtm_transactions*.

Note: All transaction data from vTMs should be sent to a specific Splunk index. This index should *only* be used for transaction data from vTMs.

- Specify the **Logs Index**. This is the index used for logs on the Splunk platform. For example, *zxtm_logs*.

Note: All log data from vTMs should be sent to a specific Splunk index. This index should *only* be used for log data from vTMs.

9. If you want Transport Layer Security (TLS) to be used during the query, select the **Query using TLS** check box.
 - You can then choose to verify the TLS connection by selecting the **Verify TLS** check box.
 - You must provide an SSL **Certificate**. To do this, either browse for the required certificate file in the **From file** property, or paste the contents of the certificate into the **From text** property.
10. Enter a **Username** and **Password** for the query authentication on the Analytics System.

FIGURE 293 Analytics: Search Endpoint Example

Add Search Endpoint

Name:

Address:

Transactions index:

Logs index:

Query using TLS: ☒

Verify TLS: ☒

Certificate: ☒ From file [Choose File](#)

☐ From text

Username:

Password:

11. (Optional) Click **Test Connection** to test the search endpoint connection using the specified properties. Success is indicated where the search endpoint can be contacted.

FIGURE 294 Analytics: Search Endpoint Test Success



If the test fails, rework your properties and re-test.

12. Click **Apply**. The new Search Endpoint resource is added to the **Search Endpoints** table.

FIGURE 295 Analytics: Added Search Endpoint


Search Endpoints

+ Add

Name	ID	Address	
JK-EP-Search-01	Search-Endpoint-6WJC-TN9I-TXB5-G4K9	demo.com:2020	<button>Test Connection</button>

13. (Optional) Expand the Search Endpoint resource entry to view its full details.
14. (Optional) Test a listed search endpoint at any time by clicking the **Test Connection** button in the **Test** column of the summary entry for the endpoint. Success is indicated where the search endpoint can be contacted.

FIGURE 296 Analytics: Search Endpoint Test Success



Note: You must also create all required Collection Endpoint resources. See [“Adding a Collection Endpoint Resource to the Services Director” on page 289](#).

Enabling Analytics on a vTM Cluster

Once all analytics resources are in place on the Services Director (see [“Creating Analytics Resources” on page 285](#)), you can enable vTM analytics on a cluster of vTMs. There are two steps to this process:

- Using the Services Director VA GUI, update each vTM in the cluster to use a Feature Pack that includes a SKU that supports vTM analytics. See [“Applying a Feature Pack to Registered Instances” on page 140](#).
- Using the Services Director VA GUI, update the vTM cluster to use an Analytics Profile. This configures all vTMs in the cluster to generate the analytics data specified by its supported Log Export Types. The vTM is automatically assigned an endpoint in the Analytics System from the pool of Collection Endpoints on the Services Director, and the single Search Endpoint resource. See [“Adding an Analytics Profile to a vTM Cluster” on page 297](#).

Once complete, all vTMs in the vTM cluster will generate analytics data and transmit this data to an assigned collection endpoint in the Analytics System. You are then able to query this data from the Services Director, see [“Working with Analytics Data on the Services Director” on page 298](#).

Adding an Analytics Profile to a vTM Cluster

To enable analytics on all vTMs in a cluster, you must apply an analytics profile to the vTM cluster.

This single action results in the automatic update of every vTM in the cluster by cluster replication, and completes the configuration of analytics from the Services Director.

Note: Before you can enable analytics in a vTM cluster, you must ensure that all vTMs in the cluster use a Feature Pack that supports analytics. See [“Applying a Feature Pack to Registered Instances” on page 140](#).

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The Home page appears.
3. Click the **Services** menu, and then click **Services Director > vTM Clusters**.

The **vTM Clusters** page appears.

- Expand the cluster that you want to update.

FIGURE 297 Analytics: View Details for a vTM Cluster

vTM Clusters

The screenshot shows the 'vTM Clusters' management interface. At the top, there is a table with columns: Cluster Name, Type, In Use, Analytics Profile, Backup Schedule, Next Backup Time, Action, Last Action, and Last Action Status. The first row shows 'Carmine-Cluster' with Type 'Discovered', In Use 'checked', Analytics Profile 'N/A', Backup Schedule 'N/A', and an 'Add' button. Below the table, there is a form to update the cluster details: Cluster Name (Carmine-Cluster), Owner (JK), Analytics Profile (N/A), Backup Schedule (N/A), and Number of Backups (5). There are 'Apply' and 'Revert' buttons. At the bottom, there is a section for backups with columns: Backup Name, Description, Date, Retain, and Actions. A message states: 'There are no backups currently available for this cluster'.

- Select the required vTM cluster and click **Apply**.

FIGURE 298 Analytics: Update a vTM Cluster

The screenshot shows the 'Add vTM Cluster' dialog box. It has a title bar with a close button. The form contains: Cluster Name (carmine_cluster), Owner (JK), Analytics Profile (None), and Backup Schedule (None). There is an 'Add' button. A dropdown menu for Backup Schedule is open, showing options: Event Only, Audit & Event, and Audit Only.

The cluster update tests all required analytics resources. See [“Creating Analytics Resources” on page 285](#) if issues arise.

If all required analytics resources are in place, the cluster updates. After this process is complete, all vTMs in the cluster are updated by cluster replication, and analytics becomes enabled on all vTMs.

Analytics data then starts to accumulate in the Analytics System, and can be queried from the Services Director Analytics interface. See [“Working with Analytics Data on the Services Director” on page 298](#).

Working with Analytics Data on the Services Director

Note: This functionality is available to all Services Director customers.

The Services Director can then use the vADC Analytics Application to query the Analytics System and present the data as a variety of analytics graphs.

- The Analytics Dashboard. This provides a fixed view onto a selection of graphs, to provide high-level information. See [“Accessing the vADC Analytics Application” on page 299](#).
- A number of individual analytics graph types. Each graph type focus on one graphical representation type. This includes:
 - *Tree graphs*. See [“Using the Sankey Diagram” on page 325](#).
 - *Table graphs*. See [“Using the Table Graph” on page 332](#).
 - *Charts*. See [“Using Charts” on page 334](#).
 - *Dataset graphs*. See [“Using the Dataset View” on page 360](#).

Each graph uses a common set of filters to limit data. These filters can be changed at any time:

- The **Data Selector**. See [“Choosing a Data Metric” on page 302](#).
- The **Time Selector**. See [“Choosing a Time Period” on page 303](#).
- The **Sampling Selector**. See [“Choosing a Sampling Ratio” on page 305](#).
- The **Component Filter**. See [“Working with the Component Filter” on page 307](#).
- The **Extended Filter**. See [“Working with the Extended Filter” on page 316](#).

Graph-specific behaviours then enable manipulation of displayed data, filtering of results, and drill-down.

- The log data saved from one or more servers. See [“Working with the Logs View” on page 363](#).

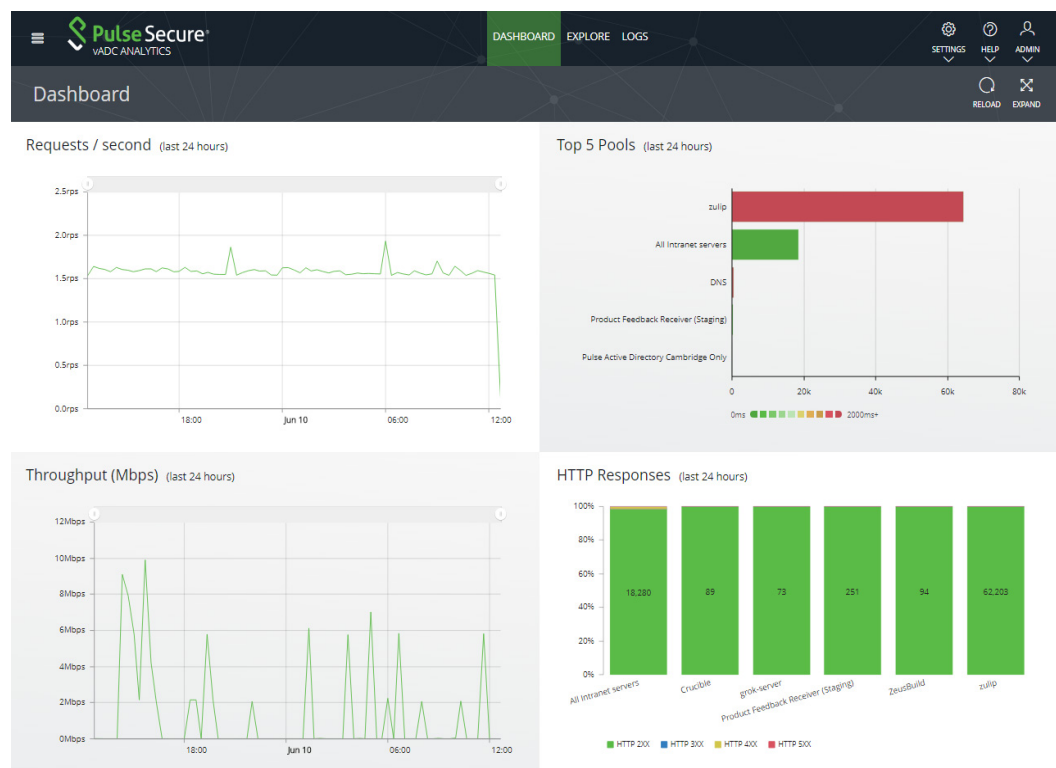
Accessing the vADC Analytics Application

The **vADC Analytics** application provides access to a dashboard and individual analytics graphs.

1. Access your Services Director VA from a browser, using its Service Endpoint IP Address.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **Analytics: Dashboard** and log into the **vADC Analytics** application using the Services Director credentials.

The **vADC Analytics** application starts in a new window, starting with the **Dashboard** page. This page presents a view onto a selection of fixed graphs within a single page. Each graph provide a high-level view of your analytics data. For example:

FIGURE 299 The Analytics Dashboard




You cannot interact with these graphs. However, you can access individual graph types to perform any required analysis.

The graph types are:

- *Tree graphs.* See [“Working with the Extended Filter”](#) on page 316.
- *Table graphs.* See [“Using the Table Graph”](#) on page 332.
- *Charts.* See [“Using Charts”](#) on page 334.
- *Dataset graphs.* See [“Starting the Dataset View”](#) on page 361.

You can return to the Dashboard at any time by clicking **Dashboard**.

FIGURE 300 Analytics: The Dashboard Button



Returning to the Services Director VA

When you are in the **vADC Analytics** application, you may want to return to the Services Director VA.

Note: When you start the **vADC Analytics** application from the Services Director VA, a separate browser tab is started. The tab for the Services Director VA may still be available.


1. In the **vADC Analytics** application, click the **Menu** button.

FIGURE 301 Analytics: The Menu Button



The menu appears.

FIGURE 302 Analytics: The Menu Button




2. Click **Go To Services Director**.

The Services Director VA appears.

Choosing a Data Metric

The **Metric Selector** is one of the standard filters that apply to all analytics graph types.

FIGURE 303 Analytics: Filtering Using Data Metric



The selected data metric limits the scope of data to a specific measurement type, such as total throughput or requests per second.

The total data for the analytics graph is defined by the combined settings from the **Time Selector**, the **Metric Selector**, the **Sampling Selector**, the **Component Filter**, and the **Extended Filter**. Any of these criteria can be changed at any time, and the analytics graph will automatically update to reflect your selections.

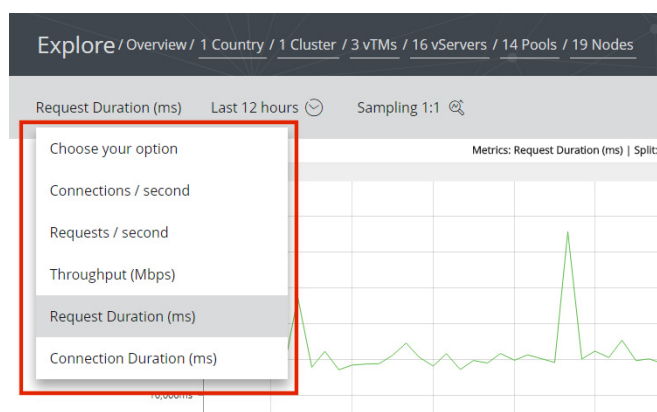
Note: Also see [“Choosing a Time Period” on page 303](#), [“Choosing a Sampling Ratio” on page 305](#), [“Working with the Component Filter” on page 307](#) and [“Working with the Extended Filter” on page 316](#).

1. Start the **vADC Analytics** application, see [“Accessing the vADC Analytics Application” on page 299](#).
2. Access the required analytics graph type.

The page for the selected graph type appears. This page includes standard filter controls as well as graph-specific controls.

3. Click the **Metric Selector** to view all available data metric options.

FIGURE 304 Analytics: Selecting a Data Metric



In this example, you can select total throughput (expressed as Megabits per second), or the number of requests per second.

Note: Some metrics do not support percentiles, and are disabled when percentiles are in use.


4. Click your required data metric.

Once your selection is made, the analytics graph updates automatically, based on the current settings for the **Time Selector**, **Metric Selector**, **Sampling Selector**, **Component Filter**, and **Extended Filter**.

Choosing a Time Period

The **Time Selector** is one of the standard filters that apply to all analytics graph types.

FIGURE 305 Analytics: Filtering Using Time Period



The selected time period limits the scope of data to a specific period of time, which typically ends at the current time. You can also select historical ranges if required.

The total data for the analytics graph is defined by the combined settings from the **Time Selector**, the **Metric Selector**, the **Sampling Selector**, the **Component Filter**, and the **Extended Filter**. Any of these criteria can be changed at any time, and the analytics graph will automatically update to reflect your selections.


Note: Also see [“Choosing a Data Metric” on page 302](#), [“Choosing a Sampling Ratio” on page 305](#), [“Working with the Component Filter” on page 307](#) and [“Working with the Extended Filter” on page 316](#).

1. Start the **vADC Analytics** application, see [“Accessing the vADC Analytics Application” on page 299](#).
2. Access the required analytics graph type.

The page for the selected graph type appears. This page includes standard filter controls as well as graph-specific controls.

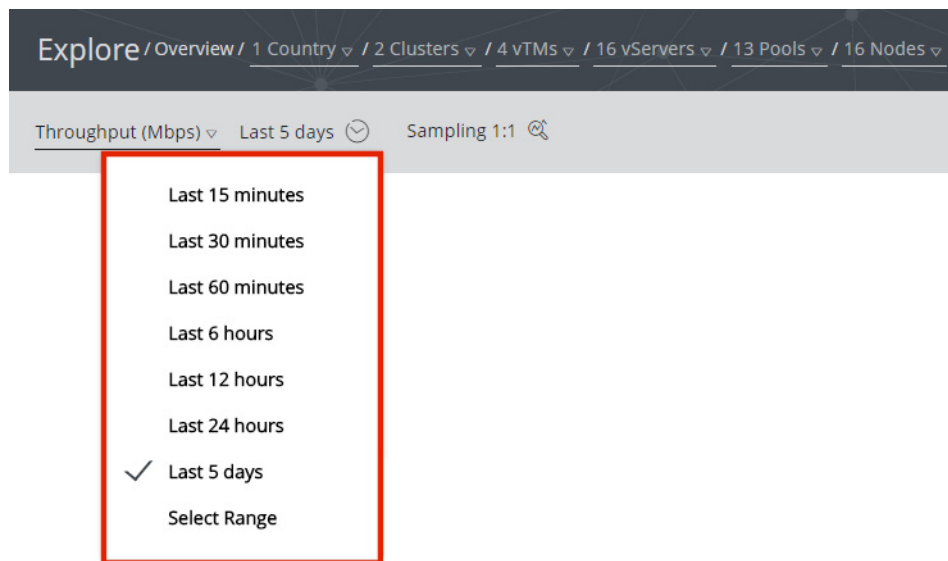
3. Click the **Time Selector** button to view the list.

FIGURE 306 Analytics: The Time Period Button




A list of fixed time periods appears.

FIGURE 307 Analytics: Selecting a Recent Time Period



4. (Optional) If you want to include the most recent data in your graph, select the time period that you require from the list. For example, to view data for the last hour, click **Last 60 minutes**.
5. (Optional) If you want to include a time period that is not specifically listed, or which does not end at the current time, click **Select Range**. The current list is replaced with a pair of filters that control the start and end of the required time period.

FIGURE 308 Analytics: Selecting a Time Period Range



Click on either filter to access standard date/time selection tools.

6. (Optional) To return to a fixed time period, click the **Time Selector** button and make the required selection.

Once your time period selection is complete, the **Component Filter** updates automatically to include only those components for which data was received during the requested period. See [“Working with the Component Filter” on page 307](#).


The analytics graph also updates automatically, based on the current settings for the **Time Selector**, **Metric Selector**, **Sampling Selector**, and **Extended Filter**.

Choosing a Sampling Ratio

The **Sampling Selector** is one of the standard filters for analytics graph types.

Note: The **Sampling Selector** does *not* apply to the Dataset View. See [“Using the Dataset View” on page 360](#).

FIGURE 309 Analytics: Filtering Using Sampling Ratio



By default, an analytics graph includes all events for its specified criteria. However, in some situations you might want to retrieve a smaller *sampled* set of events, instead of retrieving the entire event set:

- You may want to determine the nature of a large data set without processing every event.
For example, for a very large dataset where you wish to study trends, a sampled dataset will be retrieved faster and is likely to indicate all significant trends.
- You may want to perform a quick search to check that expected events are being returned from the current search criteria.

A *sampling ratio* is the probability of any single event being included in the total result set. For example, if the sample ratio value is *1:100*, each event has a 1 in 100 chance of being included in the results. The selection of each event is independent. It is possible that many events will be included from the first 100 events, or that none of these will be included.

If you to re-run a sampling search, different *specific* results will almost certainly be returned.

A range of sampling ratios from *1:10* to *1:10000* are supported in Services Director. A *1:10* sampling ratio retrieves the most data and is the most representative of source data. A sampling ratio of *1:10000* retrieves the least data and is less representative. A sampling ratio of *1:1* indicates that all data is included. That is, that there is no sampling.

Note: Pulse Secure recommends that you use a *1:1* sampling ratio (that is, there is no sampling) whenever it is practical. If sampling is required, your search should always retrieve as much data as practical. That is, if a *1:10* sampling ratio produces acceptable results, do not proceed to using a *1:100* sampling ratio.

Note: Where analytics events are used to calculate totals (such as *Throughput* and *Requests per Second*), sampling should be used with caution. All totals will be approximated for the entire dataset based on the sample, and its heading will be marked with an asterisk to indicate that all numbers are approximate. As the sampling ratio increases, the accuracy of this approximation decreases.

The total data for the analytics graph is defined by the combined settings from the **Time Selector**, the **Metric Selector**, the **Sampling Selector**, the **Component Filter**, and the **Extended Filter**. Any of these criteria can be changed at any time, and the analytics graph will automatically update to reflect your selections.

Note: Where a sampled set of results does not include a selected value for a specific **Component Filter** category, the selected value for the filter is cleared.


Also see [“Choosing a Data Metric” on page 302](#), [“Choosing a Time Period” on page 303](#), [“Working with the Component Filter” on page 307](#) and [“Working with the Extended Filter” on page 316](#).

1. Start the **vADC Analytics** application, see [“Accessing the vADC Analytics Application” on page 299](#).
2. Access the required analytics graph type.

The page for the selected graph type appears. This page includes standard filter controls as well as graph-specific controls.


3. Click the **Sampling Selector** to view all available data metric options.

FIGURE 310 Analytics: Choosing a Sampling Ratio



4. Click your required sampling ratio.

FIGURE 311 Analytics: Displayed Sampling Ratio



After you have chosen to use sampling, any data that is the result of sampling is indicated, either by:

- The column heading for the value is prefixed by an asterisk.
- The data value itself is prefixed by an asterisk.
- Any "equals" signs are replaced by "approximately equal to" signs.

Once your selection is made, the analytics graph updates automatically, based on the current settings for the **Time Selector**, **Metric Selector**, **Sampling Selector**, **Component Filter**, and **Extended Filter**.


Note: Also see [“Choosing a Data Metric” on page 302](#), [“Choosing a Time Period” on page 303](#), [“Working with the Component Filter” on page 307](#) and [“Working with the Extended Filter” on page 316](#).

Working with the Component Filter

The **Component Filter** is one of the standard filters that apply to all analytics graph types.

Note: There is also an extended set of filters, see [“Working with the Extended Filter” on page 316](#).

FIGURE 312 Analytics: The Component Filter



The total data for the analytics graph is defined by the combined settings from the **Time Selector**, the **Metric Selector**, the **Sampling Selector**, the **Component Filter**, and the **Extended Filter**. Any of these criteria can be changed at any time, and the analytics graph will automatically update to reflect your selections.

Note: Also see [“Choosing a Data Metric” on page 302](#), [“Choosing a Time Period” on page 303](#), [“Choosing a Sampling Ratio” on page 305](#), and [“Working with the Extended Filter” on page 316](#).

Understanding the Component Filter

The **Component Filter** has six component categories (**Location**, **Clusters**, **vTMs**, **vServers**, **Pools** and **Nodes**). You can make selections in all, some or no categories as required.

Note: The **Location** category can be configured to be based on *Continents*, *Countries* or *Cities*, see [“Configuring the Location Category” on page 310](#).

The **Component Filter** only lists components for which analytics data is recorded, restricted by:


- The current **Time Selector** setting. See [“Choosing a Time Period” on page 303](#).
- The current **Sampling Selector** setting. See [“Choosing a Sampling Ratio” on page 305](#).
- Any selections already made in the **Component Filter**.
- Any selections made in the **Expanded Filter**. See [“Working with the Extended Filter” on page 316](#).

When you make a selection, the **Component Filter** categories can update automatically:

- Where no data is recorded for an individual component after any restrictions (or selections) are applied, the component is omitted from its component category list.
- If you make a selection for a component category, the **Component Filter** displays and highlights the selection. All other categories (both higher-level and lower-level) for which no selection is made may be updated so that only entries that relate to the most recent selection are listed.
- If no component selection is made for a component category, the current number of components for the category is displayed.

All selections are highlighted:


FIGURE 313 Analytics: Highlighted Selections



You can clear a single component type selection by expanding its list and clicking **Reset Filter**.


You can completely reset the **Component Filter** at any time by clicking the **Reset** button:

FIGURE 314 Analytics: Resetting Component Filters




You can refresh retrieved analytics data by clicking the **Reload** button. For example, to refresh the analytics data for the *Last 6 hours*:

FIGURE 315 Analytics: Reloading Analytics Data




You can configure an extended set of filters in addition to the **Component Filter** by clicking the **Filter** button, see [“Working with the Extended Filter” on page 316](#).

FIGURE 316 Analytics: Accessing the Extended Filter



You can maximize the space within the browser by clicking the **Expand** toggle.

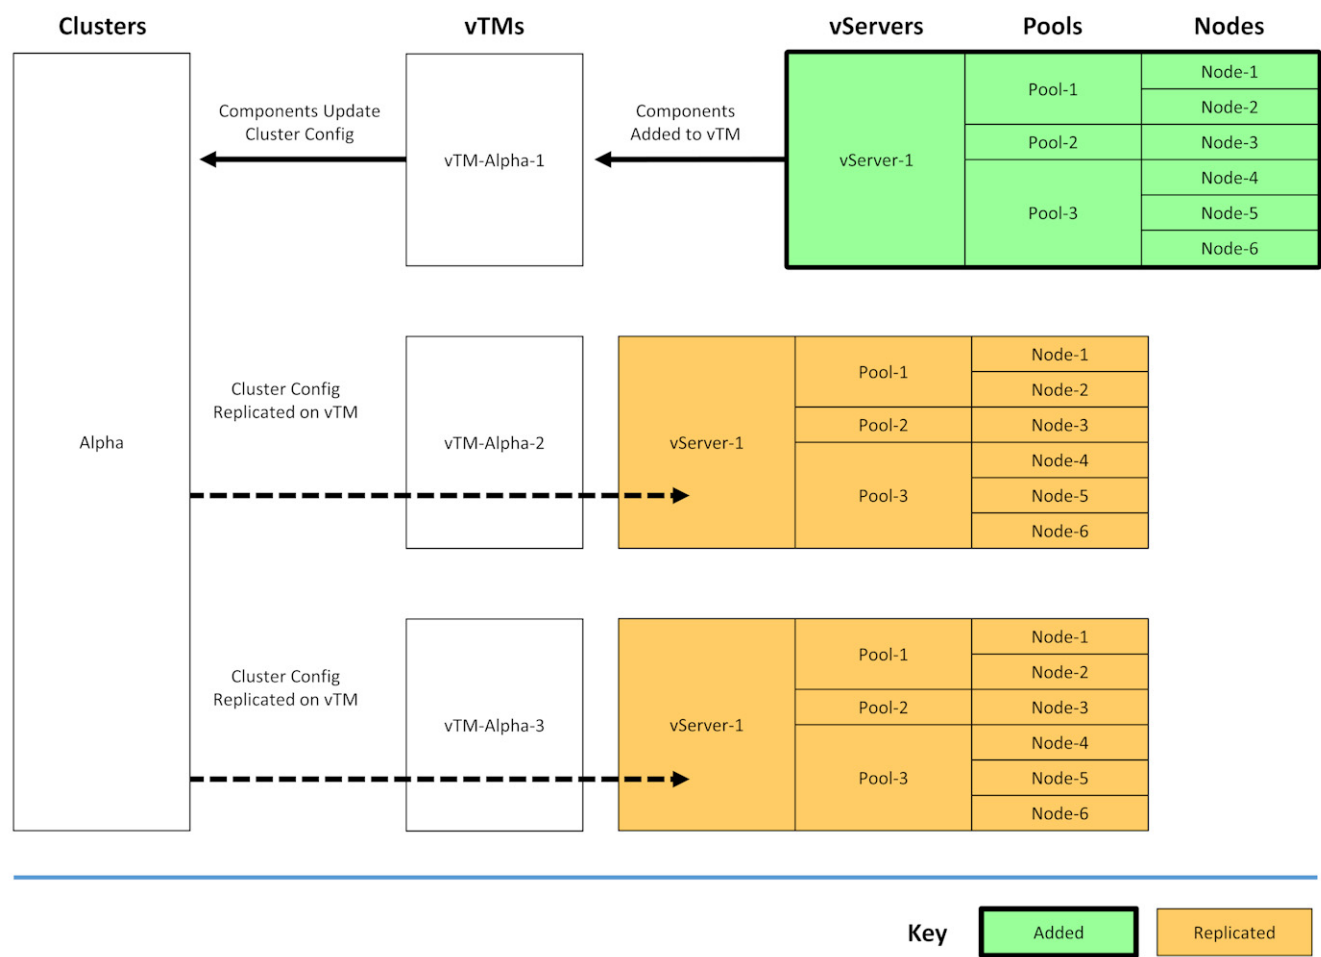
FIGURE 317 Analytics: Maximizing the Data Area for the Browser



Understanding Cluster-Level Replication of Components

The configuration of vServers, Pools and Nodes is a cluster-level operation. That is, the configuration of vServers, Pools and Nodes on any vTM is automatically duplicated on all other vTMs in the cluster, using cluster replication. The names and configurations of these resources will be identical.

FIGURE 318 Analytics: Cluster-Level Replication of Components



In larger clusters, this will result in large numbers of identically named components within the cluster. To address this issue, all duplicate names are eliminated in the **Component Filter**. See [“Understanding Component Filter Categories” on page 310](#).

Understanding Component Filter Categories

The **Component Filter** has six component categories.

- **Location** category. This category enables you to filter by the geographic location (where known), and can be configured to be based on *Continents*, *Countries* or *Cities*, see [“Configuring the Location Category” on page 310](#).
- **Clusters** category. Each vTM can be a member of one cluster only, but multiple clusters may be visible from the Services Director. You can make a single cluster selection if required.
- **vTMs** category. This lists all vTMs within the selected Cluster, or for all listed Clusters if no Cluster is selected. You can make a single vTM selection if required.
- **vServers** category. This lists all vServers within the selected vTM, or for all listed vTMs if no vTM is selected. You can make a single vServer selection if required.
- **Pools** category. This lists all Pools within the selected vServer, or for all vServers if no vServer is selected. You can make a single pool selection if required.
- **Nodes** category. This lists all back-end Nodes within the selected Pool, or for all Pools if no Pool is selected. You can make a single pool selection if required.

Listed components in all categories are restricted automatically by all previous category selections, and by selections to other filters. Only components for which analytics data exists after all selections and filters are applied are included.

Note: All categories can include an entry listed as “None”. This can indicate, for example:

- Incomplete transaction data. That is, a transaction that starts but does not complete, such as might occur during equipment failure.
- Data was retrieved from a cache rather than by forwarding the request.

Note: Cluster-Level configurations such as vServers, pools and nodes will result in repeated component names across all vTMs in a cluster. Component names are not repeated within a category list. A single selected component can refer to many actual components, which can be further explored by making additional selections. See [“Understanding Cluster-Level Replication of Components” on page 309](#).

Configuring the Location Category

The Location category enables you to filter by the geographic location of the remote client IP address (where this can be determined). The geographic location can be based on *Continents*, *Countries* or *Cities*.

Where the geographic location of a remote client IP address cannot be identified, such as in a private network, the data is added to a generic Location category grouping called *<Unknown>*.


Data from the following standard private networks (as defined by the Internet Assigned Numbers Authority) can be included as a named Location category grouping.

- *10.0.0.0/8*. This represents the reserved address for 24-bit subnetworks (class A network).

- 172.16.0.0/12. This represents the reserved address for 20-bit subnetworks (class B network).
- 192.168.0.0/16. This represents the reserved address for 16-bit subnetworks (class C network).

When any of these options are selected, their network can appear in the Location category of the **Component Filter**. For example:


FIGURE 319 Analytics: Location Category Includes Subnetwork



Configuring the Location Category

1. Click **Settings** on the toolbar to access the analytics settings.

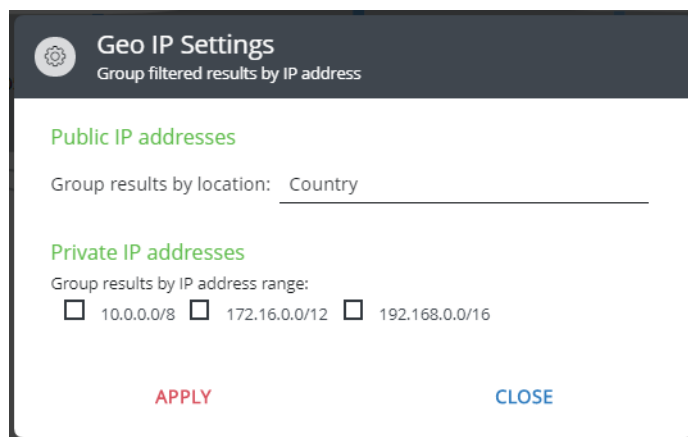
FIGURE 320 Analytics: Accessing the Analytics Settings



2. On the pull-down menu, click **Geo IP Settings**.

The **Geo IP Settings** dialogue box appears.

FIGURE 321 Analytics: The Settings Dialog Box



3. Under **Public IP addresses**, select the required geographical grouping. That is, *Continents*, *Countries* or *Cities*.
4. Under **Private IP addresses**, select any required standard private networks. That is, *10.0.0.0/8*, *172.16.0.0/12* or *192.168.0.0/16*.
5. Click **Apply**.

Example 1: Hierarchic Selection

When you use the **Component Filter** as a hierarchy, you make left-to-right selections to narrow the scope of a graph to specific components. For example:

FIGURE 322 Analytics: Hierarchic Selection: Hierarchy of Components

Geographic	Clusters	vTMs	vServers	Pools	Nodes
Asia	Alpha	vTM-Alpha-1	vServer-1	Pool-1	Node-1
				Pool-2	Node-2
				Pool-3	Node-3
		vTM-Alpha-2	vServer-1	Pool-1	Node-4
				Pool-2	Node-5
				Pool-3	Node-6
Europe	Beta	vTM-Beta-1	vServer-1	Pool-1	Node-1
				Pool-2	Node-2
				Pool-3	Node-3
		vTM-Beta-2	vServer-1	Pool-4	Node-4
				Pool-1	Node-5
				Pool-2	Node-6
		vTM-Beta-1	vServer-1	Pool-1	Node-1
				Pool-2	Node-2
				Pool-3	Node-3
		vTM-Beta-2	vServer-1	Pool-4	Node-4
				Pool-1	Node-5
				Pool-2	Node-6

Key

Component

Required Path

In this example, analytics data exists for all end-to-end paths shown, taking into account the selected time range (see **“Choosing a Time Period” on page 303**). The required end-to-end path is marked in green; data that was created for this path is required for an analytics graph.

To deliver the required information to the graph, you can use the **Component Filter** to select the components on the path, one at a time, working left-to-right. The listed options adjust automatically as each selection is made.

For this example:

1. Expand each category in turn and examine the lists. Components for all possible paths are shown:
 - There are two continents in the **Location** category.
 - There are two clusters, each of which is in a separate continent.
 - There are four vTMs across the two clusters.
 - There is one listed vServer. There are four vServers in total across the four vTMs, but there is a single repeating name because of cluster replication. All duplicates are removed. See [“Understanding Cluster-Level Replication of Components” on page 309](#).
 - There are four pools. There are fourteen pools in total across the four vTMs, but there are repeating names because of cluster replication. All duplicates are removed.
 - There are six nodes. There are 24 nodes in total across the four vTMs, but there are repeating names because of cluster replication. All duplicates are removed.
2. Expand the **Location** category. Two continents are listed: *Asia* and *Europe*. Select *Asia*,
3. Expand the **Clusters** category. Two clusters are listed: *Alpha* and *Beta*. Select *Alpha*.
4. Expand the **vTMs** category. Only the two vTMs in the Alpha Cluster are listed: *vTM-Alpha-1* and *vTM-Alpha-2*. Select *vTM-Alpha-2*.
5. Expand the **vServers** category. Only *vServer-1* is listed, as this is the only vServer in the selected vTM. Select *vServer-1*.
6. Expand the **Pools** category. Three pools are listed, as these are the pools within the selected vTM: *Pool-1*, *Pool-2* and *Pool-3*. Select *Pool-3*.
7. Expand the **Nodes** category. Three nodes are listed, as these are the nodes within the selected vTM: *Node-4*, *Node-5* and *Node-6*. Select *Node-5*.

All selections are now complete. The analytics graph will use all data for the pathway between the *Asia* continent and *Node-5* on *vTM-Alpha-2*. This represents an end-to-end connection.

Note: The analytics graph updates after every selection.

Note: You can also reach the same result using a different number of **Component Filter** selections, using a flexible selection approach. See [“Example 2: Flexible Component Selection” on page 314](#).

Example 2: Flexible Component Selection

When you use the **Component Filter** to explore analytics data, you can select from any component category at any time, subject to restrictions placed by previous selections.

For example, here is a possible hierarchy of components:

FIGURE 323 Analytics: Flexible Selection: Hierarchy of Components

Geographic		Clusters		vTMs		vServers		Pools		Nodes
Asia		Alpha		vTM-Alpha-1		vServer-1		Pool-1		Node-1
								Node-2		
								Pool-2		Node-3
								Pool-3		Node-4
										Node-5
										Node-6
				vTM-Alpha-2		vServer-1		Pool-1		Node-1
								Node-2		
								Pool-2		Node-3
								Pool-3		Node-4
										Node-5
										Node-6
Europe		Beta		vTM-Beta-1		vServer-1		Pool-1		Node-1
								Node-2		
								Pool-2		Node-3
								Pool-3		Node-4
								Pool-4		Node-5
										Node-6
				vTM-Beta-2		vServer-1		Pool-1		Node-1
								Node-2		
								Pool-2		Node-3
								Pool-3		Node-4
								Pool-4		Node-5
										Node-6

In this example, analytics data exists for all end-to-end paths shown, taking into account the selected time range (see **“Choosing a Time Period” on page 303**).

You can explore the analytics data, and view the filtered results, by making selections in any category. For this example:

- Expand each category in turn and examine the lists. Components for all possible paths are shown:
 - There are two continents in the **Location** category.
 - There are two clusters, each of which is in a separate continent.

- There are four vTMs across the two clusters.
 - There is one vServer. There are four vServers in total across the four vTMs, but there is a single repeating name because of cluster replication. All duplicates are removed. See [“Understanding Cluster-Level Replication of Components” on page 309](#).
 - There are four pools. There are fourteen pools in total across the four vTMs, but there are repeating names because of cluster replication. All duplicates are removed.
 - There are six nodes. There are 24 nodes in total across the four vTMs, but there are repeating names because of cluster replication. All duplicates are removed.
2. Expand the **Nodes** category. Six nodes are listed: *Node-1*, *Node-2*, *Node-3*, *Node-4*, *Node-5* and *Node-6*. Select *Node-5*. This selection includes all Nodes called *Node-5*, of which there are four. (see below)
 3. Expand the **vTMs** category. Four vTMs are listed, as each of these vTMs contains a Node called *Node-5*. Select *vTM-Alpha-2*.

The two selections have now identified a single pathway between the *Asia* continent and *Node-5* on *vTM-Alpha-2*. This represents an end-to-end connection.

FIGURE 324 Analytics: Flexible Selection: Implicit Path Based on Two Selections

Geographic		Clusters		vTMs		vServers		Pools		Nodes			
Asia		Alpha		vTM-Alpha-1		vServer-1		Pool-1		Node-1			
								Node-2					
								Pool-2		Node-3			
								Pool-3		Node-4			
										Node-5			
										Node-6			
				Pool-1		Node-1							
						Node-2							
						Pool-2		Node-3					
								Node-4					
Pool-3	Node-5												
	Node-6												
	Europe		Beta			vTM-Beta-1			vServer-1		Pool-1		Node-1
											Node-2		
Pool-2											Node-3		
Pool-3											Node-4		
Pool-4											Node-5		
											Node-6		
						Pool-1			Node-1				
									Node-2				
Pool-2									Node-3				
									Node-4				
	Pool-3	Node-5											
		Node-6											
Pool-4		Node-1											
		Node-2											
	Pool-1	Node-3											
		Node-4											
Pool-2		Node-5											
		Node-6											

Key

Component

Selected Component

Implicit Path

Key

Component

Selected Component

Implicit Path

No more selections are supported without clearing one of the category selections.

Note: The analytics graph updates after each selection.

Note: You can also reach the same result using a different number of **Component Filter** selections, using an hierarchic selection approach. See [“Example 1: Hierarchic Selection” on page 312](#).

Working with the Extended Filter

The **Extended Filter** is one of the standard filters that apply to all analytics graph types.

When used, one or more clauses appear in the **Extended Filter**. All of these must be satisfied for a data item to be included in any analytics graph. For example:

FIGURE 325 Analytics: Extended Filter Panel

Type to filter options...	Choose an operator	Select a value
HTTP Response Code	IS	400
HTTP Response Code	IS	500
Transaction Duration	GREATER THAN	1000

CANCEL APPLY FILTER

The use of the **Extended Filter** is described in the following sections:

- [“Starting the Extended Filter” on page 318](#)
- [“Adding Clauses to the Extended Filter” on page 318](#)
- [“Understanding Extended Filter Clauses” on page 321](#)
- [“Understanding Implicit Logical Operators Between Clauses” on page 322](#)

Note: If you create an **Extended Filter** clause that is based on one of the standard **Component Filter** categories, the available values for that category will also be restricted in the **Component Filter**.


The total data for the analytics graph is defined by the combined settings from the **Time Selector**, the **Metric Selector**, the **Sampling Selector**, the **Component Filter**, and the **Extended Filter**. Any of these criteria can be changed at any time, and the analytics graph will automatically update to reflect your selections.

Note: Also see [“Choosing a Data Metric” on page 302](#), [“Choosing a Time Period” on page 303](#), [“Choosing a Sampling Ratio” on page 305](#), and [“Working with the Component Filter” on page 307](#).

Starting the Extended Filter


To start the **Extended Filter**, click the **Filter** toggle on the toolbar.

FIGURE 326 Analytics: Accessing the Extended Filter



The **Extended Filter** appears at the bottom of the browser window. When it is started for the first time, it contains no clauses.

FIGURE 327 Analytics: Extended Filter Panel



To minimize the extended filter, click the **Filter** toggle again.

Adding Clauses to the Extended Filter

To add one or more clauses to the **Extended Filter**, perform the following steps.


1. Start the **Extended Filter**, see [“Starting the Extended Filter” on page 318](#).

The **Extended Filter** appears at the bottom of the browser window.

2. In the **Extended Filter**, either:
 - Type the name of the required filter option (field) for the clause, OR
 - Expand the list of filter options (fields) and select the required option for the clause.

For example:

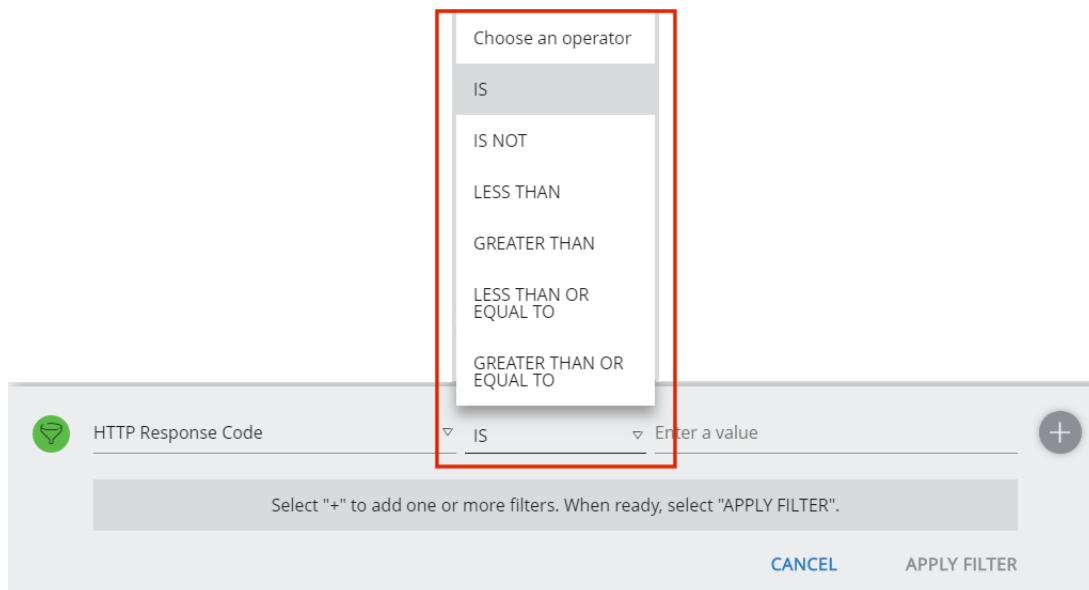
FIGURE 328 Analytics: Selecting an Option for a Clause



See [“Understanding Extended Filter Clauses” on page 321](#) for details of clauses.

- Expand the list of operators, and select the required operator for the clause. For example:

FIGURE 329 Analytics: Selecting an Operator for a Clause

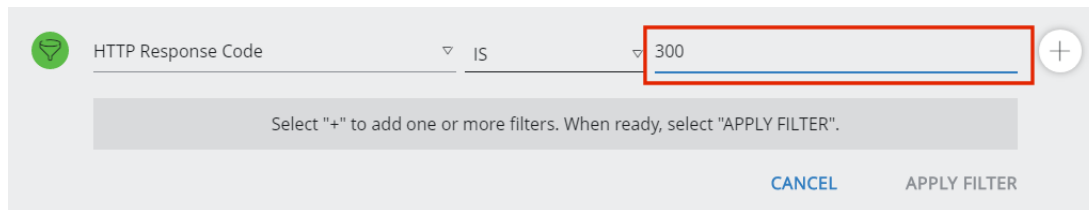


The screenshot shows the Analytics interface with a filter clause being configured. The dropdown menu for operators is open, displaying the following options: IS, IS NOT, LESS THAN, GREATER THAN, LESS THAN OR EQUAL TO, and GREATER THAN OR EQUAL TO. The 'IS' operator is currently selected. The filter clause is 'HTTP Response Code' with the operator 'IS' and a placeholder 'Enter a value'. Below the clause, there is a text prompt: 'Select "+" to add one or more filters. When ready, select "APPLY FILTER"'. At the bottom right, there are 'CANCEL' and 'APPLY FILTER' buttons.

Note: This list is tailored to the selected filter option.

- Type the required search value for the clause. For example:

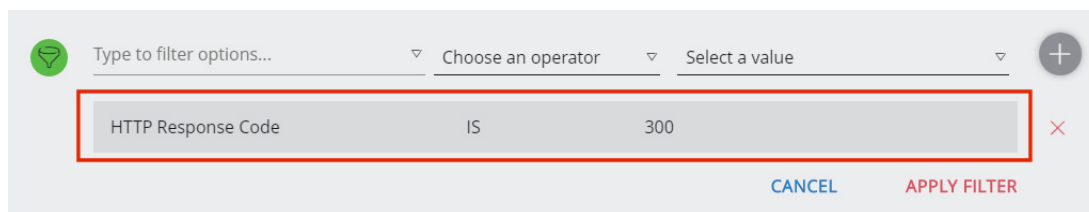
FIGURE 330 Analytics: Selecting an Value for a Clause



The screenshot shows the Analytics interface with the filter clause 'HTTP Response Code' and operator 'IS'. The value '300' has been entered into the text input field. Below the clause, there is a text prompt: 'Select "+" to add one or more filters. When ready, select "APPLY FILTER"'. At the bottom right, there are 'CANCEL' and 'APPLY FILTER' buttons.

- Click the + button. The clause is added to the list of clauses. For example:

FIGURE 331 Analytics: Adding the First Clause



The screenshot shows the Analytics interface with the first clause added to the list. The clause is 'HTTP Response Code' with the operator 'IS' and the value '300'. Below the clause, there is a text prompt: 'Select "+" to add one or more filters. When ready, select "APPLY FILTER"'. At the bottom right, there are 'CANCEL' and 'APPLY FILTER' buttons.

- Repeat steps 2 to 5 to add more clauses. For example:

FIGURE 332 Analytics: Adding the Additional Clauses

Type to filter options...	Choose an operator	Select a value	
HTTP Response Code	IS	300	×
OR			
HTTP Response Code	IS	400	×
Transaction Duration	GREATER THAN	1000	×

CANCEL APPLY FILTER

Note: Implicit logical operators are applied automatically to the list of clauses, see [“Understanding Implicit Logical Operators Between Clauses” on page 322](#).

Note: The **Extended Filter** does not display the word “AND”. All listed clauses after the first are related with an AND unless an OR is displayed.

- Click **Apply** to apply all listed clauses to the current analytics graph type.
- (Optional) To minimize the extended filter at any time, click the **Filter** toggle. When the **Extended Filter** is populated with one or more clauses, it minimizes to the bottom of the browser window and remains visible. For example:

FIGURE 333 Analytics: Extended Filter Panel Minimized

(HTTP Response Code IS 400 × OR HTTP Response Code IS 500 ×) Transaction Duration IS 1000 ×

Understanding Extended Filter Clauses

The **Extended Filter** is specified as a list of user-defined clauses. Each clause identifies:

- A field in the transaction data that was exported by a vTM to the analytics repository.
- A condition that relates to the field.
- A value for the condition.

That is:

```
<field> <condition> <value>
```

For example:

```
Remote Client Port IS 123
```

The supported conditions and values for a clause depend upon the specified field:

- Numeric fields can support one of more of the following conditions:
 - *IS*. For example: Remote Client Port IS 8080
 - *IS NOT*. For example: Remote Client Port IS NOT 8100
 - *LESS THAN*. For example: Transaction Duration LESS THAN 30
 - *LESS THAN OR EQUAL TO*. For example: Transaction Duration LESS THAN OR EQUAL TO 17
 - *GREATER THAN*. For example: Transaction Duration GREATER THAN 23
 - *GREATER THAN OR EQUAL TO*. For example: Transaction Duration GREATER THAN OR EQUAL TO 40
 - *IS PRESENT*. For example: Transaction Duration IS PRESENT
 - *IS ABSENT*. For example: Transaction Duration IS ABSENT
- String fields support the following conditions:
 - *IS*. For example: Protocol IS "HTTP"
 - *IS NOT*. For example: Protocol IS NOT "FTP"
 - *CONTAINS*. For example: Protocol CONTAINS "TP"
 - *DOES NOT CONTAIN*. For example: Protocol DOES NOT CONTAIN "FT"
 - *IS PRESENT*. For example: Protocol IS PRESENT
 - *IS ABSENT*. For example: Protocol IS ABSENT

- Boolean fields support the following conditions:
 - *IS*. For example: HTTP Response Server Keep Alive IS TRUE
 - *IS PRESENT*. For example: HTTP Response Server Keep Alive IS PRESENT
 - *IS ABSENT*. For example: HTTP Response Server Keep Alive IS ABSENT

The user does not define the logical relationships *between* the various clauses using *explicit logical operators*,

Rather, the **Extended Filter** is subject to *implicit logical operators* that are imposed automatically by the **vADC Analytics Application**, see [“Understanding Implicit Logical Operators Between Clauses” on page 322](#).

Understanding Implicit Logical Operators Between Clauses

The user can define one or more **Extended Filter** clauses to manage the information that is included in analytics graphs. See [“Understanding Extended Filter Clauses” on page 321](#).

The user does not define the logical relationships between extended filter clauses using *explicit logical operators*. Rather, the **Extended Filter** clauses are subject to *implicit logical operators* that are imposed automatically by the *vADC Analytics Application*.

- All clauses that reference a *single field* using “IS” or “CONTAINS” operator are automatically related via an implicit **OR** logical operator. For example, the following clauses reference the same field:

```
Field-X IS 10
Field-X IS 20
Field-X IS 50
```

This is equivalent to:

```
Field-X IS 10
OR Field-X IS 20
OR Field-X IS 50
```

- All other clauses are automatically related via an implicit **AND** logical operator. For example:

```
Field-A GREATER THAN 10
Field-A LESS THAN OR EQUAL TO 20
Field-B IS NOT "Halo"
Field-C IS "CBG"
Field-D IS NOT 66
Field-E IS PRESENT
```

This is equivalent to:

```
Field-A GREATER THAN 10
AND Field-A LESS THAN OR EQUAL TO 20
AND Field-B IS NOT "Halo"
AND Field-C IS "CBG"
AND Field-D IS NOT 66
AND Field-E IS PRESENT
```

- A list of clauses can combine both of these clause types:


```
Field-X IS 10
Field-X IS 20
Field-A GREATER THAN 10
Field-A LESS THAN OR EQUAL TO 20
Field-B IS NOT "Halo"
Field-C IS "CBG"
Field-D IS NOT 66
Field-E IS PRESENT
Field X IS 50
```

This is equivalent to (with **OR** terms grouped together):

```
(Field-X IS 10
OR Field-X IS 20
OR Field-X IS 50)
AND Field-A GREATER THAN 10
AND Field-A LESS THAN OR EQUAL TO 20
AND Field-B IS NOT "Halo"
AND Field-C IS "CBG"
AND Field-D IS NOT 66
AND Field-E IS PRESENT
```

In all cases, the resulting extended filter is applied to the analytics graph.

The **Extended Filter** does not display the word "AND". All listed clauses after the first are related with an AND unless an OR is displayed. For example:

FIGURE 334 Analytics: Multiple Extended Filter Clauses

Type to filter options...	Choose an operator	Select a value	
HTTP Response Code	IS	300	×
OR			
HTTP Response Code	IS	400	×
Transaction Duration	GREATER THAN	1000	×
HTTP Response Header Content-Type	IS PRESENT		×
HTTP Response Header Content-Encoding	IS PRESENT		×

CANCEL APPLY FILTER

In this example, the clauses are related as follows:

```
(HTTP Response Code IS 300 OR HTTP Response Code IS 400)
AND Transaction Duration GREATER THAN 1000
AND HTTP Response Header Content-Type IS PRESENT
AND HTTP Response Header Content-Encoding IS PRESENT
```

When the **Extended Filter** is minimized in the browser window, the clauses appear as follows:

FIGURE 335 Analytics: Multiple Extended Filter Clauses

(HTTP Response Code IS 300 × OR HTTP Response Code IS 400 ×) Transaction Duration GREATER THAN 1000 ×
 HTTP Response Header Content-Type IS PRESENT × HTTP Response Header Content-Encoding IS PRESENT ×


Using the Sankey Diagram

The supported tree graph is a Sankey diagram. This is a specific type of flow diagram, in which the width of the graph lines is proportional to the flow quantity between each pair of points.

For analytics purposes, the width of the line on the Sankey diagram shows proportional flow of the chosen data metric (see [“Choosing a Data Metric” on page 302](#)).

Flow is calculated for all end-to-end connections between the geographic areas and nodes in your vTM cluster, and displayed according to included components. For example:

FIGURE 336 Tree Graphs: Sankey Diagram



To display a Sankey diagram, see [“Starting the Sankey Diagram” on page 326](#).

Once a Sankey diagram is displayed, you can focus on your analytics data as follows:

- [“Selecting Included Components for your Sankey Diagram” on page 326](#).
- [“Focusing on a Component in a Sankey Diagram” on page 329](#).
- [“Focusing on a Path in a Sankey Diagram” on page 331](#).

You can also update the following controls at any time:


- The **Component Filter**, see [“Working with the Component Filter” on page 307](#).
- The **Metric Selector**, see [“Choosing a Data Metric” on page 302](#).
- The **Time Selector**, see [“Choosing a Time Period” on page 303](#).

The scope of the Sankey diagram updates immediately to include end-to-end connections that meet all selection criteria.

Starting the Sankey Diagram


1. Start the **vADC Analytics** application, see [“Accessing the vADC Analytics Application” on page 299](#).
2. Click **Explore** to access individual analytics graphs.

FIGURE 337 Accessing Analytics Graphs




Alternatively, click the **Menu** button, and then click **Explorer**.

FIGURE 338 Analytics: The Menu Button



3. Finally, click the **Tree** graph type.

FIGURE 339 Starting the Dataset View



The required graph type appears.

Selecting Included Components for your Sankey Diagram

By default, the Sankey diagram includes all six component categories:

- Location. This can be configured to be based on *Continents*, *Countries* or *Cities*, see [“Configuring the Location Category” on page 310](#).

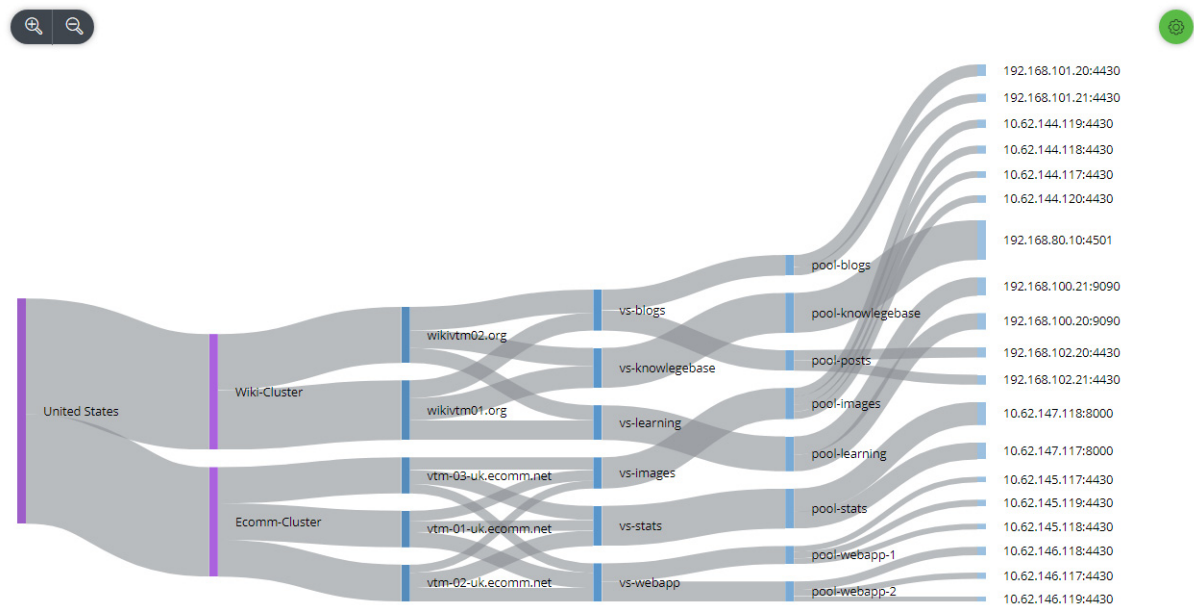
Note: Where data events are collected for more than ten country/city locations, each location is ranked according to the number of data events collected. The top ten locations are displayed individually in the Sankey diagram, and all locations after the tenth are displayed as a single entry named "Rest of the World".

- Clusters
- vTMs
- vServers
- Pools
- Nodes

You can exclude specific component types from the diagram if required.

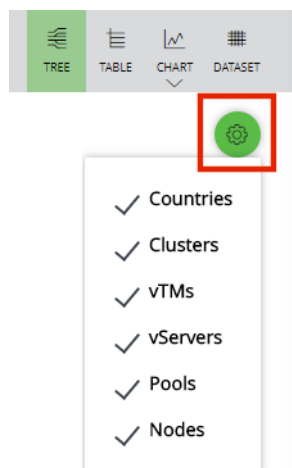
1. Display a Sankey diagram. See [“Starting the Sankey Diagram” on page 326](#). For example:

FIGURE 340 Analytics: Sankey Diagram




2. Click the **Settings** button to display a check list of component types. For example:

FIGURE 341 Analytics: Sankey Diagram Component Settings



Note: In this example, the **Location** category is set to the *Countries* setting. This can also be set to *Continent* or *City*, see [“Configuring the Location Category” on page 310](#).

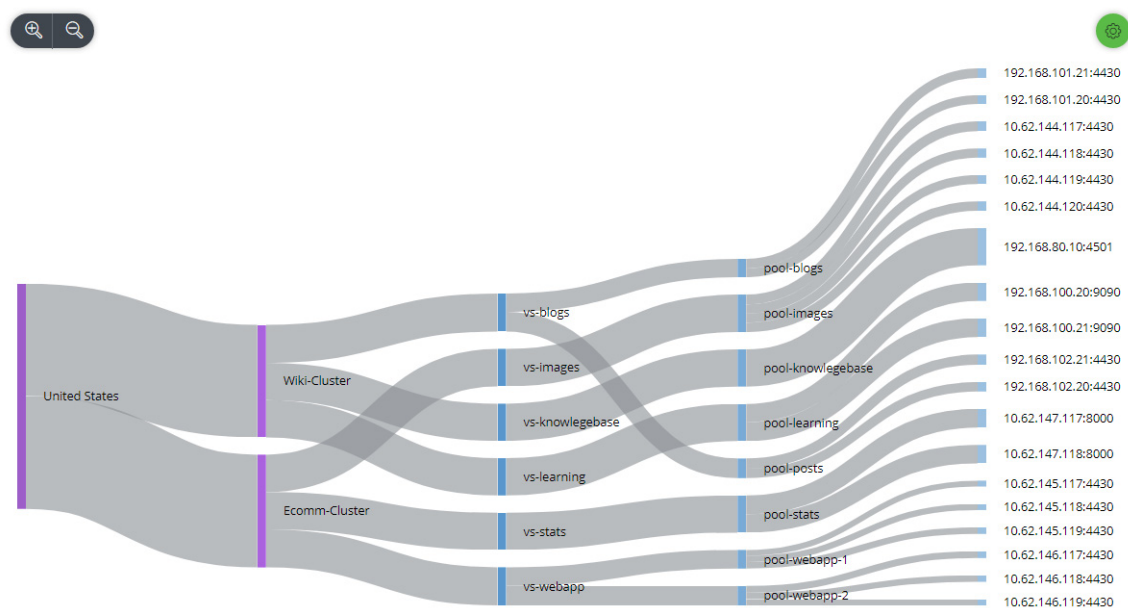
FIGURE 342 Analytics: Sankey Diagram Alternate Location Component Settings



3. Select a component type to include/exclude it.


For example, after excluding vTMs:

FIGURE 343 Analytics: Sankey Diagram Excluding vTMs



For example, after excluding both vServers and pools:

FIGURE 344 Analytics: Sankey Diagram Excluding vServers and Pools




Focusing on a Component in a Sankey Diagram

You can focus on a specific component in the Sankey diagram, which updates the graph to include only those end-to-end connections that include the selected component.


1. Display a Sankey diagram. See [“Starting the Sankey Diagram” on page 326](#). For example:

FIGURE 345 Analytics: Sankey Diagram




2. In the Sankey diagram, hover the mouse pointer over the required component to display:
 - An indication of all end-to-end paths passing through the node.
 - The name of the node. For example:

FIGURE 346 Analytics: Viewing Details for a Sankey Component



3. Click the node. The Sankey diagram updates to include all end-to-end connections that include the selected component. For example:

FIGURE 347 Analytics: Focusing on a Sankey Component




Note: You can also focus on a specific path in the Sankey diagram, see [“Focusing on a Path in a Sankey Diagram” on page 331](#).

Focusing on a Path in a Sankey Diagram

You can focus on a single path in the Sankey diagram, which updates the graph to include only those end-to-end connections that include the selected node.


1. Display a Sankey diagram. See [“Starting the Sankey Diagram” on page 326](#). For example:

FIGURE 348 Analytics: Sankey Diagram



2. In the Sankey diagram, hover the mouse pointer over the required path to see its details. For example:

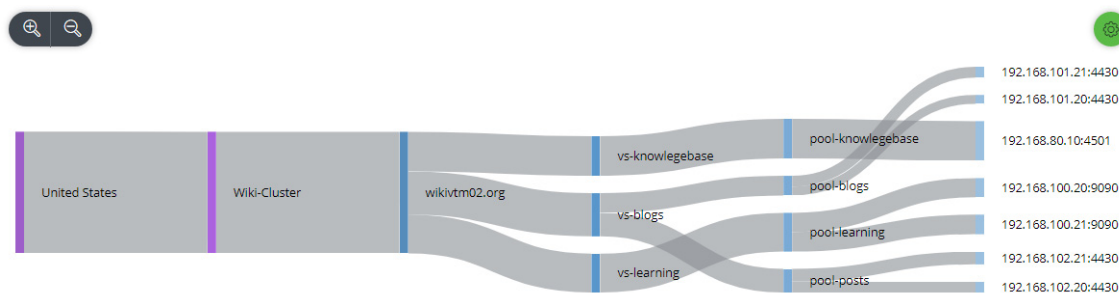
FIGURE 349 Analytics: Viewing Details for a Sankey Path



Note: When sampling is applied to the dataset, this is indicated by an asterisk prefix on the heading. For example, **Throughput** is replaced by ***Throughput**.

- Click the path. The Sankey diagram updates to include all end-to-end paths that include the selected path. For example:

FIGURE 350 Analytics: Focusing on a Sankey Path

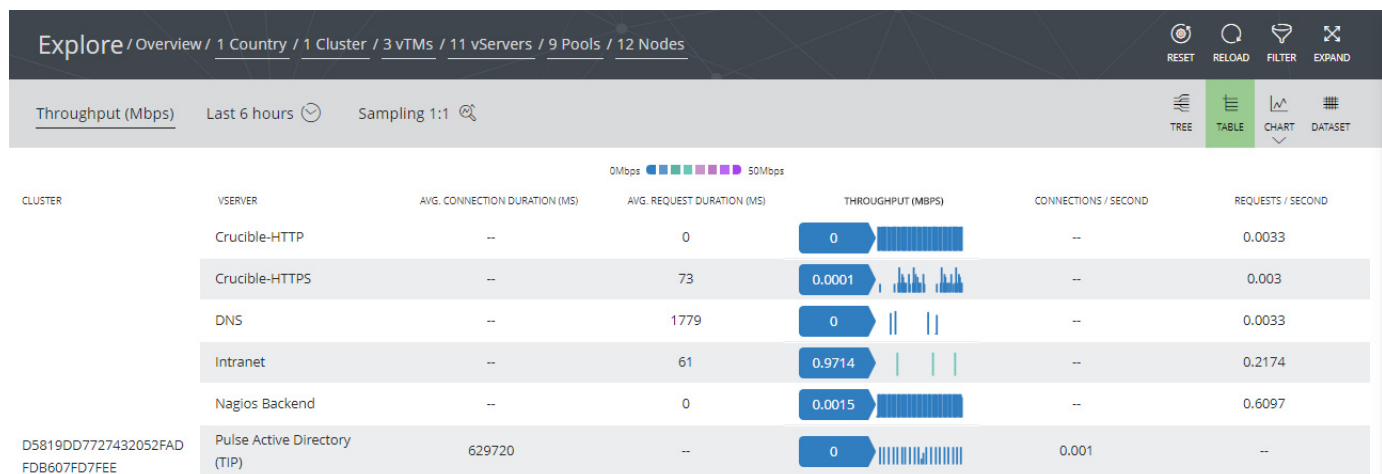


Note: You can also focus on a specific component in the Sankey diagram, see [“Focusing on a Component in a Sankey Diagram” on page 329](#).

Using the Table Graph

The supported Table Graph is a per-vServer summary of all of the available metrics. The graph also includes a sparkline that shows trends in the currently data for all selected criteria. For example:

FIGURE 351 Table Graph



To display a Table Graph, see [“Using Charts” on page 334](#).

You can also update the following controls at any time:


- The **Component Filter**, see [“Working with the Component Filter” on page 307](#).
- The **Metric Selector**, see [“Choosing a Data Metric” on page 302](#).
- The **Time Selector**, see [“Choosing a Time Period” on page 303](#).

The scope of the Table Graph updates to include and-to-end connections that meet all selection criteria.

Starting the Table Graph


1. Start the **vADC Analytics** application, see [“Accessing the vADC Analytics Application” on page 299](#).
2. Click **Explore** to access individual analytics graphs.

FIGURE 352 Accessing Analytics Graphs



3. Finally, click the **Table** graph type.

FIGURE 353 Starting the Dataset View



The required graph type appears.

Understanding the Table Graph


The Table Graph can include the following measurements:

- Cluster
- vServer
- Average Connection Duration (milliseconds). This property contains a connection duration measurement for a protocol such as TCP.
- Average Request Duration (milliseconds). This property contains a request duration measurement for a protocol such as HTTP or HTTPS.
- Throughput (Mbits per second)
- Connections per Second.
- Requests per Second.

Some of these measurements will be blank, depending on the protocol in use, and on the selected data metric, see [“Choosing a Data Metric” on page 302](#).


Where sampling is used, this is indicated by an asterisk prefix in the column headings. For example:

FIGURE 354 Table Graph: Sampled Output



The measurement that matches your selected data metric (see [“Choosing a Data Metric” on page 302](#)) is supplemented with a “sparkline” graphic. This graphic visually summarizes measurements across the required time range, with an overall colour coding. For example:

FIGURE 355 Table Graph: Sparkline



Using Charts

The Primary Chart displays values for the current data metric over time. Optionally, this can be split by component type.


A set of secondary graphs on tabs underneath the Primary Chart provide deeper analysis and comparisons with the main chart. These are:

- The **Comparative Analysis** tab, see [“Performing Comparative Analysis” on page 350](#).
- The **Alternative Views** tab, see [“Viewing the Horseshoe Diagram” on page 355](#).
- The **HTTP Response Codes** tab, see [“Viewing HTTP Response Codes” on page 358](#).
- The **Top Events** tab, see [“Viewing Top Events” on page 359](#).

Starting the Chart


1. Start the **vADC Analytics** application, see [“Accessing the vADC Analytics Application” on page 299](#).
2. Click **Explore** to access individual analytics graphs.

FIGURE 356 Accessing Analytics Graphs



3. Click the **Chart** graph type.

FIGURE 357 Starting the Dataset View



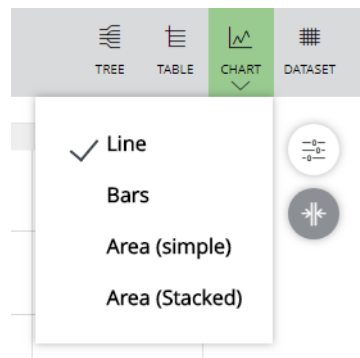
4. Select the required graph type, see [“Chart Types” on page 336](#).

The required graph type appears.

Chart Types


There are four chart types supported, each of which is accessed from the **Chart** pull-down menu.

FIGURE 358 Charts: Line Chart Types




- Line charts. For example:

FIGURE 359 Charts: Line chart




Line charts support splits. For example, if split by vTM:

FIGURE 360 Charts: Line chart split by vTM




- Bar charts. For example:

FIGURE 361 Charts: Bar chart



Bar charts support splits. For example, if split by vTM:


FIGURE 362 Charts: Line chart split by vTM



Note: When splits are used, bar charts are presented as stacked data.


- Simple area charts. For example:

FIGURE 363 Charts: Area chart



Area charts support splits. For example, if split by vTM:


FIGURE 364 Charts: Line chart split by vTM



- Stacked area charts. This chart type requires split data, as different data sets are cumulatively stacked vertically.

For example:

FIGURE 365 Charts: Stacked area chart split by vTM



Using a Logarithmic Vertical Axis


A logarithmic scale is a nonlinear scale that is used when there is a large range of quantities.

If an axis uses a logarithmic scale, each displayed value is ten times bigger than the one beneath it, as it is based on orders of magnitude; large values become closer together visually, and more differentiation is possible for values that are closer to zero.

Linear Scales and Logarithmic Scales

The following diagrams compare the same data displayed using linear and logarithmic scales.


FIGURE 366 Charts: Linear Scale



In this example:

- The vertical axis is marked from 0Mbps to 40Mbps in linear 10Mbps increments.
- The smaller values (many less than 1Mbps) are hard to read (and to differentiate from zero/missing), because of the huge difference between them and the larger values on the linear scale.

FIGURE 367 Charts: Logarithmic Scale



In this example:

- The vertical axis is marked from 0.01Mbps to 100Mbps, with each value ten times bigger than the last:
 - 0.01Mbps
 - 0.1Mbps
 - 1Mbps
 - 10Mbps
 - 100Mbps


- The smaller values are easier to read, because the logarithmic scale is more detailed at that level.

Assigning a Linear Scale or Logarithmic Axis Scale

To select the required axis scale:

1. Click the **Settings** button.

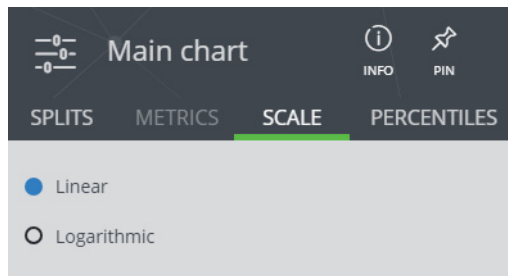
FIGURE 368 Charts: Settings Button



2. On the menu, select **Scale**.

The **Main Chart** settings panel appears with the **Scale** tab selected.

FIGURE 369 Main Charts: Scale Tab



3. (Optional) click **Pin** to fix the panel to the side of the main display. This remains until unpinned.
4. Select the required axis scale, either:
 - *Linear*
 - *Logarithmic*
5. The Main Chart updates automatically.

Viewing Percentile Values

You can view percentile values within the main chart.

Note: Percentiles are disabled when splits are in use, see [“Splitting the Primary Chart” on page 345](#).


Note: Some data metrics do not support percentiles. These metrics are disabled when percentiles are in use, see [“Viewing Percentile Values” on page 340](#).

When you view percentiles, the main data line is replaced by three customizable percentile lines. By default, these lines are:

- The 99th percentile.
- The 95th percentile.
- The 50th percentile.

For example:


FIGURE 370 Percentiles in the Main Chart



To replace the main data line by between one and three percentile lines on the main chart:


1. View the main chart. For example:

FIGURE 371 Main Chart



2. Click **Settings** for the main chart.

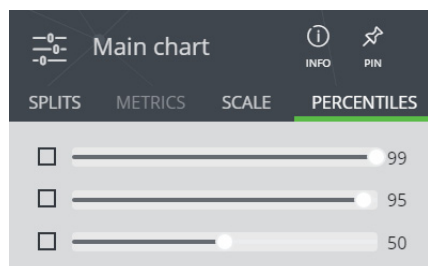
FIGURE 372 Charts: Settings Button



3. Select a chart metric that supports percentiles. That is, either:
 - *Request Duration (ms)*
 - *Connection Duration (ms)*
4. In the menu, select **Percentile**.


The **Main Chart** settings panel appears with the **Percentiles** tab selected.

FIGURE 373 Main Chart: Percentiles Tab




5. (Optional) click **Pin** to fix the panel to the side of the main display. This remains until unpinned.
6. Select the required number of percentiles.

FIGURE 374 Main Chart: Percentiles Enabled



7. (Optional) Update the individual values of the enabled percentiles to a value between 1 and 100.
- The main chart updates automatically.


FIGURE 375 Percentiles in the Main Chart



Working with the Primary Chart

The Primary Chart displays metrics over time. For example:


FIGURE 376 Charts: The Primary Chart



Note: Where sampling is used, this is indicated by a smoothed curve.


To examine data values for a point in time, hover the mouse pointer over a line.

FIGURE 377 Charts: Examining Data Values



Note: Where sampling is used, this is indicated by an “approximately equal to” symbol, and an asterisk prefix for the value. For example:


FIGURE 378 Table Graph: Unsampled Data vs Sampled Data



Splitting the Primary Chart

Optionally, you can split the Primary Chart by component type. For example, If you split by vServer, each vServer has its own colour-coded line:

FIGURE 379 Charts: The Primary Chart Split by vServer



Note: Where there are potentially more than ten lines, only the first ten are displayed individually. The data events from all remaining lines are aggregated as a single line named "Other".


Note: When splits are used, bar charts are presented as stacked data.

Note: When splits are used, percentiles are disabled. See [“Viewing Percentile Values” on page 340](#).

To split the Primary Chart by a selected criteria:

1. Click the **Settings** button.

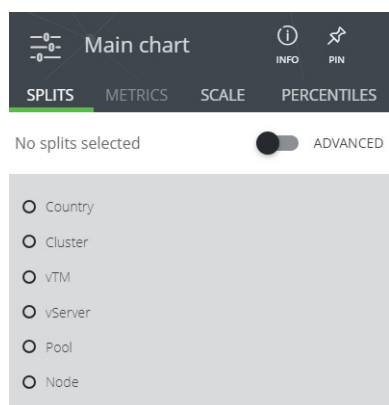
FIGURE 380 Charts: Settings Button



2. In the menu, select **Splits**.

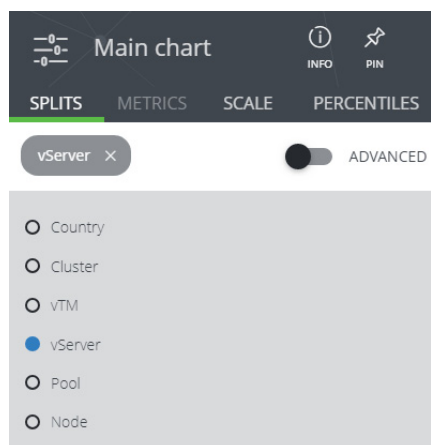
The **Main Chart** panel appears.

FIGURE 381 Charts: Main Chart Settings Panel



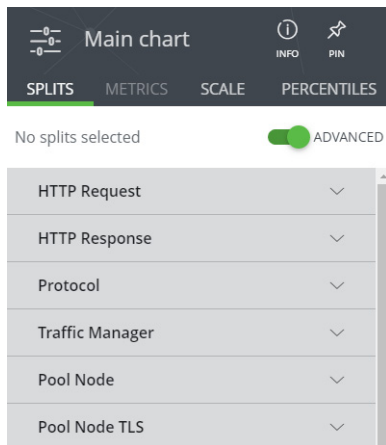
3. (Optional) click **Pin** to fix the panel to the side of the main display. This remains until unpinned.
4. Then, choose a split category. Either:
 - If you want to split the Primary Chart using one of the basic component categories, select the **Basic** switch setting, and then select the required category. For example, **vServer**.

FIGURE 382 Charts: Basic Split



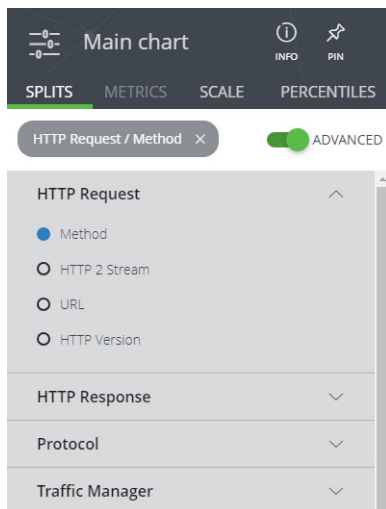
- If you want to split the Primary Chart using more specific criteria, select the **Advanced** switch setting.

FIGURE 383 Charts: Advanced Chart Split Selection



Then, locate and expand the required category, and select the required criteria. For example:

FIGURE 384 Charts: Making an Advanced Chart Split Selection




In both cases, once a selection is applied, the Primary Chart updates to reflect the selection.

Note: Where there are potentially more than ten lines, only the first ten are displayed individually. The data events from all remaining lines are aggregated as a single line named "Other".

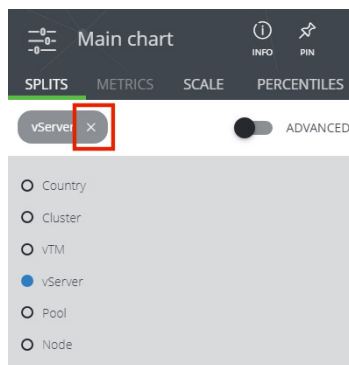
- To examine data values for a point in time, hover the mouse pointer over the split lines. For example:

FIGURE 385 Charts: Examining Data Values in a Split Chart



- (Optional) To temporarily remove a split line from the display, click on its legend entry to the left of the graph. The line is then removed, and the graph is re-drawn. Click the legend again to re-include the line.
- (Optional) To return to an un-split Primary Chart, delete the current selection on the **Main Chart** panel.

FIGURE 386 Charts: Clearing a Split Selection




Focusing on a Time Range on the Primary Chart

You can focus the Primary Chart to a specific time range in the graph.


- Display the Primary Chart (split if required). For example:

FIGURE 387 Charts: Primary Chart




2. Drag across a time range in the graph. For example:

FIGURE 388 Charts: Selecting a Time Range




The graph updates to temporarily focus on the selected time range. The displayed section (a proportion within the original graph) is indicated by the sliders above the graph. For example:

FIGURE 389 Charts: Focused Graph



3. (Optional) Click the **Focus** button to permanently update the selected time range of the graph.


FIGURE 390 Charts: Permanently Re-Focusing the Graph



The position of each slider also updates.

4. (Optional) Click the Show All button to return the graph to its original time range.

FIGURE 391 Charts: Returning a Graph to Original Time Range




The position of each slider also updates.

Performing Comparative Analysis

The Comparative Analysis tab enables you to view two different data metrics in a separate graph. This graph is based on the Primary Chart (see [“Starting the Chart” on page 334](#)). For example:

FIGURE 392 Analytics: Comparative Analysis Graph



Note: Control of the display settings for the Comparative Analysis graph is similar to that used on the main chart. However, splits and percentiles can only be applied when the comparative view contains a single data metric.

Creating a Comparative Analysis Graph


1. Display the Primary Chart, see [“Starting the Chart” on page 334](#).

Note: Do **not** split the Primary Chart. This is not supported by the Comparative Analysis graph.

2. Select the required time period for the Primary Chart, see [“Choosing a Time Period” on page 303](#).
3. Select the required data metric for the Primary Chart, see [“Choosing a Data Metric” on page 302](#).
4. (Optional) Set the **Component Filter** to include the required components, see [“Working with the Component Filter” on page 307](#).
5. (Optional) Set the **Extended Filter** to include the required components, see [“Working with the Extended Filter” on page 316](#).


- Click the **Comparative Analysis** tab beneath the Primary Chart. The chart displays two charts, each based on a single default metric.

FIGURE 393 Analytics: Comparative Analysis Tab



- (Optional) To change the displayed metrics, click the **Settings** button in the **Comparative Analysis** tab.


FIGURE 394 Analytics: Comparative Analysis Settings



- In the menu, select **Metrics**.
The **Comp. Analysis** settings panel appears.

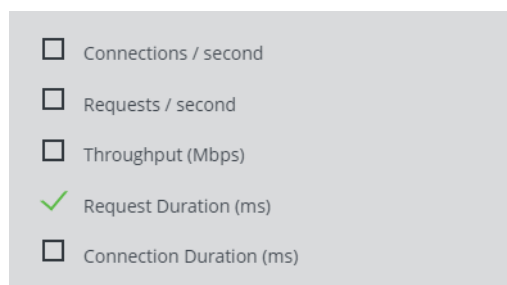
9. Click the **Metrics** tab selected. The two default metrics are indicated:

FIGURE 395 Analytics: Available Data Metrics



10. (Optional) click **Pin** to fix the panel to the side of the main display. This remains until unpinned.
11. (Optional) To switch one displayed metric for another, click the tick for a displayed metric.

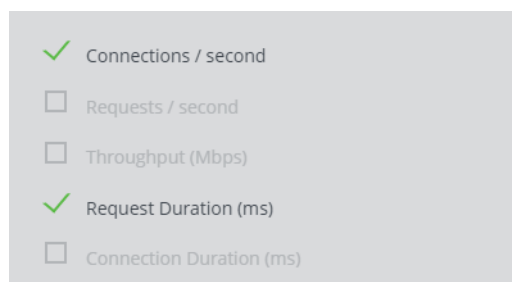
FIGURE 396 Analytics: Selecting a Data Metric



Note: Optionally, when you have a single data metric displayed in the Comparative Analysis graph, you can split the metric. You can also replace the data line with percentiles.


12. (Optional) Click the check box for the required second metric.

FIGURE 397 Analytics: Selecting a Data Metric



The Comparative Analysis graph updates.


FIGURE 398 Analytics: Updating Data Metrics



In this example, the *Connections / Second* data metric has been added. The data axis for this second metric is shown to the right of the Comparative Analysis Graph.


13. (Optional) Hover the mouse pointer over a data point in either graph to examine values in both graphs.

FIGURE 399 Charts: Examining Data Values in Two Graphs



14. (Optional) Drag the mouse pointer over either graph to temporarily re-focus both graphs.


FIGURE 400 Charts: Dragging in Either Graph



See also **“Focusing on a Time Range on the Primary Chart”** on page 348.

The graph updates to reflect the change.

FIGURE 401 Charts: Refocusing a Comparative Analysis Graph

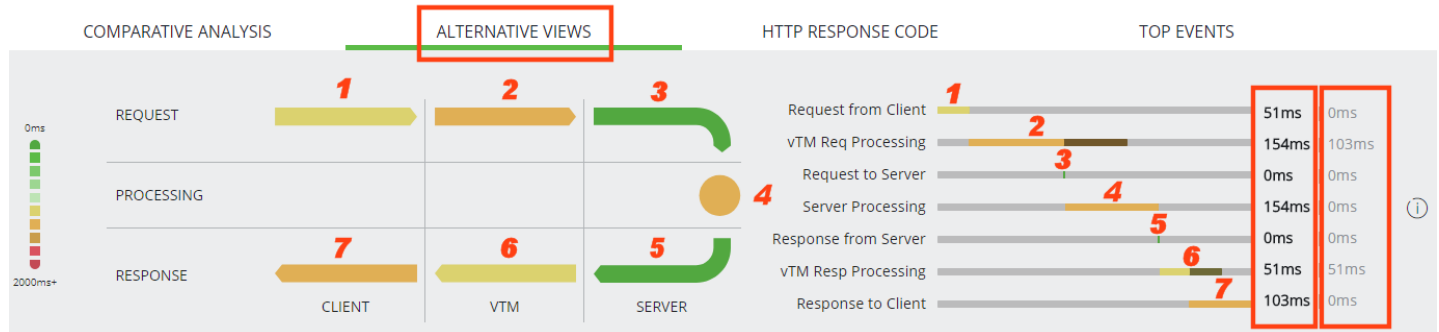


The behaviour of this focused view is the same as that described in **“Focusing on a Time Range on the Primary Chart”** on page 348.

Viewing the Horseshoe Diagram

The Horseshoe Diagram displays average timings for various activities along the receive/transmit path for client requests, based on a single vServer. Colour coding is used. For example:

FIGURE 402 Analytics: Horseshoe Diagram



Note: In this diagram, the numbers and boxes are superimposed. Descriptions are below.

The seven stages of the horseshoe diagram are:

1. **Request from Client:** The average time (in milliseconds) between the start and end of the client request reception on the vTM.
2. **vTM Req Processing:** The average time (in milliseconds) between the start of processing of the client request by the vTM, and the vTM being ready to communicate with the server. This time includes any TrafficScript processing that is required.
3. **Request to Server:** The average time (in milliseconds) between the start and end of the request being sent to the server for processing.
4. **Server Processing:** The average time (in milliseconds) for processing of the request by the server.
5. **Response from Server:** The average time (in milliseconds) between the start and end of the request being returned from the server.
6. **vTM Resp Processing:** The average time (in milliseconds) between the start of processing of the client response by the vTM. This time includes any TrafficScript processing that is required.
7. **Response to Client:** The average time (in milliseconds) between the start and end of the client response transmission from the vTM.

Next to the horseshoe diagram is a Gantt chart of timings. For each of the seven stages:

- The **Timeline** timing is for the part of the process that must complete before the vTM can begin processing the next stage. In generic Gantt chart terms, it indicates the *critical path*, and the colour associated with it is used for the matching section on the horseshoe diagram.

Note: This timing is also displayed numerically in the first column to the right of the Gantt chart.

- The **Overlap** timing is for the remainder of a process after the next process starts. For example, HTTP client requests have both a request header and a request body, but vTM request processing can begin as soon as the request header is received. As such, the two processes overlap. In generic Gantt chart terms, it indicates a *non-critical path*, and (where present), it is coloured in a darker shade of the colour used for the Timeline timing. For example, see stage 2 and 3, above.

Note: This timing is also displayed numerically in the second column to the right of the Gantt chart.

Creating a Horseshoe Diagram

1. Display the Primary Chart, see [“Starting the Chart” on page 334](#).
2. Click the **Alternative Views** tab beneath the Primary Chart. For example:


FIGURE 403 Chart View: Alternative Views



Note: The **Alternative Views** tab requires a single selected vServer.

3. (Optional) Split the Primary Chart by vServer, see [“Splitting the Primary Chart” on page 345](#). This enables you to view Charts for each vServer. For example:

FIGURE 404 Chart View: Split Primary Chart




4. Identify a single vServer using one of the following methods:

- Select the required vServer in the **Component Filter**, see [“Working with the Component Filter” on page 307](#). OR
- Identify a single vServer using an **Extended Filter** clause, see [“Working with the Extended Filter” on page 316](#). OR
- Hover the mouse pointer over the vServer lines in the Primary Chart. Then, select the required vServer by clicking on one of its data points.

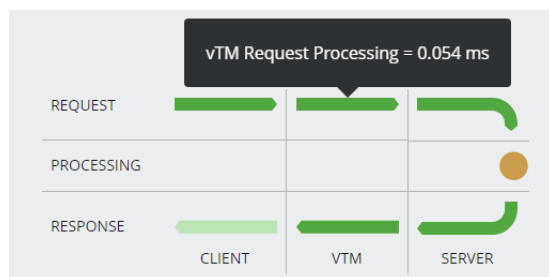
After performing one of these methods, the **Alternative Views** tab updates to show the Horseshoe Diagram for the identified vServer. For example:

FIGURE 405 Dataset View: Horseshoe Diagram for a Selected vServer




5. (Optional) Hover the mouse pointer over a section of the Horseshoe Diagram to see its value. For example:

FIGURE 406 Dataset View: Viewing a Horseshoe Diagram Value



Note: Where sampling is used, this is indicated by an asterisk prefix and an "approximately equal to" symbol. For example:

FIGURE 407 Dataset View: Sampled Horseshoe Diagram Value




- (Optional) To clear the selected vServer, expand the list of vServers in the **Component Filter**, and click **Reset filter**. See [“Working with the Component Filter” on page 307](#).

Viewing HTTP Response Codes

The HTTP Response Codes tab displays a bar chart that shows the HTTP Response codes received by the vTM pools present in the current Primary Chart. The response codes are percentage-based, and grouped into ranges of 100. For example:

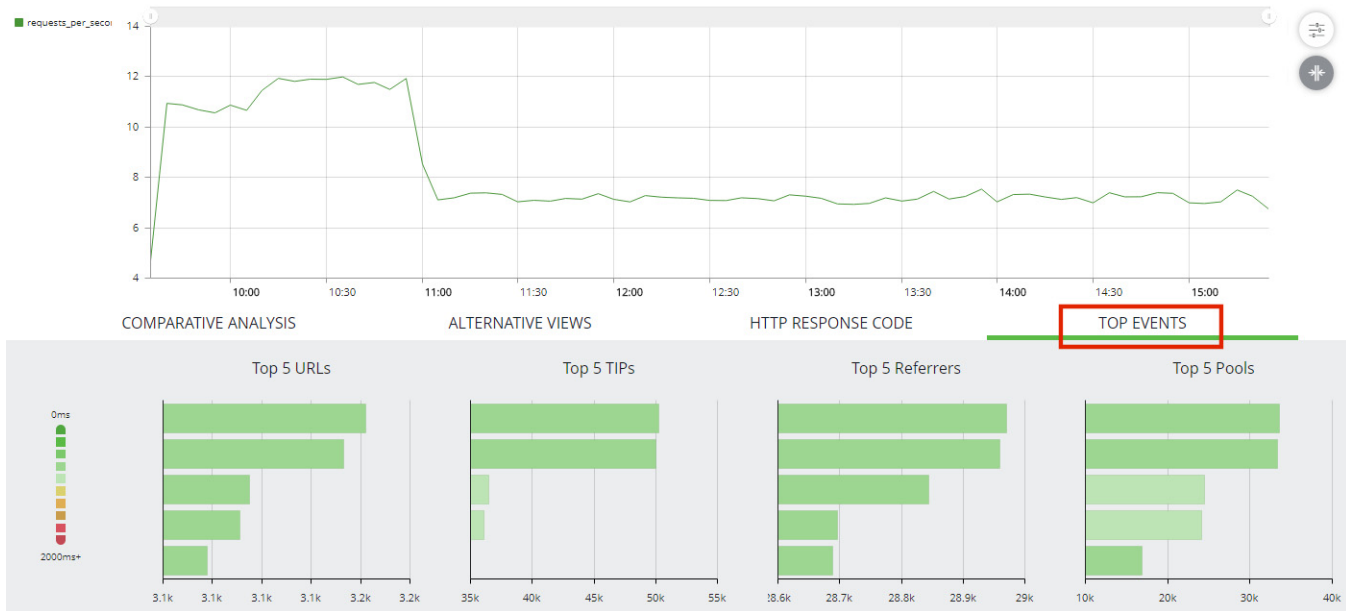
FIGURE 408 Analytics: HTTP Response Codes



Viewing Top Events

The Top Events tab displays stacked bar charts that shows HTTP Response codes for the vTM pools. The response codes are grouped into ranges of 100. For example:

FIGURE 409 Analytics: Top Events




The displayed Top Event Graphs are:

- Top 5 URLs.
- Top 5 TIPS.
- Top 5 Referrers.
- Top 5 Pools.

Hover the mouse pointer over any bar to view its details. For example:

FIGURE 410 Analytics: Viewing Top Event Details



Note: When the Primary Chart is split, the bar charts are updated to results from the split category instead of the default pools.

Note: When sampling is applied to the dataset, the entries and order of the entries in these graphs may vary between enquiries.

Using the Dataset View

The Dataset View displays the retrieved analytics data as individual rows of a table. For example:

FIGURE 411 Analytics Graphs: Dataset View

TIME	VTM	VSERVER	POOL	CLIENT IP	VIA (ADDRESS)
2018-06-10 16:31:44	intranet-2	Nagios Backend	None	10.62.164.145	10.62.130.3
2018-06-10 16:31:44	intranet-2	zulip	zulip	172.22.8.71	10.62.130.9
2018-06-10 16:31:42	intranet-2	zulip	zulip	172.22.8.103	10.62.130.9
2018-06-10 16:31:41	intranet-2	Intranet	All Intranet servers	10.62.166.176	10.62.130.3
2018-06-10 16:31:41	intranet-1	Intranet	All Intranet servers	10.62.166.176	10.62.130.4
2018-06-10 16:31:40	intranet-2	zulip	zulip	172.22.8.145	10.62.130.9
2018-06-10 16:31:39	intranet-2	zulip	zulip	10.62.160.86	10.62.130.9
2018-06-10 16:31:38	intranet-1	Nagios Backend	None	10.62.164.146	10.62.130.4
2018-06-10 16:31:38	intranet-2	Nagios Backend	None	10.62.164.146	10.62.130.3
2018-06-10 16:31:38	intranet-1	Nagios Backend	None	10.62.164.144	10.62.130.4
2018-06-10 16:31:38	intranet-2	Nagios Backend	None	10.62.164.144	10.62.130.3
2018-06-10 16:31:37	intranet-2	zulip	zulip	172.22.8.145	10.62.130.9
2018-06-10 16:31:37	intranet-2	zulip	zulip	172.22.8.156	10.62.130.9
2018-06-10 16:31:37	intranet-2	zulip	zulip	10.62.160.86	10.62.130.9

Note: Sampling is never applied to the Dataset View.

The following properties are included for all data metrics:


- Time
- vTM
- vServer
- Pool
- Client IP
- Via
- Protocol
- Node
- Duration (ms)
- Bytes In
- Bytes Out
- Completion code
- HTTP method
- HTTP code

- HTTP URL

Starting the Dataset View


1. Start the **vADC Analytics** application, see [“Accessing the vADC Analytics Application” on page 299](#).
2. Click **Explore** to access individual analytics graphs.

FIGURE 412 Accessing Analytics Graphs



3. Finally, click the **Dataset** graph type.

FIGURE 413 Starting the Dataset View



The Dataset View appears.

Viewing the Data for a Specific Row

You can view the underlying data that was used to create a specific row of the Dataset View.

1. Display the Dataset View, see [“Starting the Dataset View” on page 361](#).
2. Locate and select the required row by clicking anywhere in the row.

The **Display in Window** button for the row activates (blue).

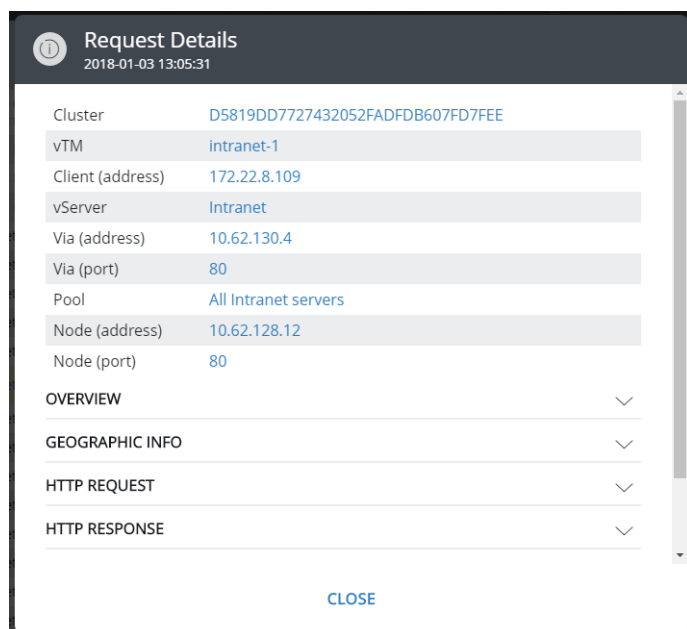
3. Click the **Display in Window** button.

FIGURE 414 Dataset View: Displaying Data for a Selected Row

Throughput (Mbps) ▾ Last 24 hours ☺ Sampling 1:1 🔍				
TIME ⓘ	VTM ⓘ	VSERVER ⓘ	POOL ⓘ	
🔍 2018-01-03 12:58:53	intranet-0	Intranet	None	
🔍 2018-01-03 13:00:47	intranet-2	Intranet	All Intranet servers	
🔍 2018-01-03 13:03:53	intranet-1	Intranet	None	
🔍 2018-01-03 13:05:31	intranet-1	Intranet	All Intranet servers	
🔍 2018-01-03 13:05:31	intranet-1	Intranet	None	

The **Request Details** window appears. This includes identifying information from the selected row.

FIGURE 415 Accessing Analytics Graphs



4. (Optional) Expand any of the sections to see the underlying data for the section:
 - Overview
 - Geographic Info
 - HTTP Request
 - HTTP Response
 - Request Trace and Timeline
 - Raw Data. This section includes entries that can be expanded to see deeper data.

Working with the Logs View

The Logs View displays retrieved log entries as individual rows of a table. For example:

FIGURE 416 Analytics Graphs: Logs View

DATE	TIME	HOST	SOURCE	SEVERITY	MESSAGE
2018-01-10	17:05:02	intranet-0	authlog	N/A	Jan 10 17:05:02 intranet-0 CRON[32224]: pam_unix(cron:session): sessi
2018-01-10	17:05:02	intranet-0	authlog	N/A	Jan 10 17:05:02 intranet-0 CRON[32224]: pam_unix(cron:session): sessi
2018-01-10	17:05:02	intranet-0	syslog	N/A	Jan 10 17:05:02 intranet-0 CRON[32225]: (root) CMD (command -v debi
2018-01-10	17:05:01	intranet-2	authlog	N/A	Jan 10 17:05:01 intranet-2 CRON[22996]: pam_unix(cron:session): sessi
2018-01-10	17:05:01	intranet-2	authlog	N/A	Jan 10 17:05:01 intranet-2 CRON[22996]: pam_unix(cron:session): sessi
2018-01-10	17:05:01	intranet-2	syslog	N/A	Jan 10 17:05:01 intranet-2 CRON[22997]: (root) CMD (command -v debi
2018-01-10	17:05:01	intranet-1	authlog	N/A	Jan 10 17:05:01 intranet-1 CRON[4507]: pam_unix(cron:session): sessio
2018-01-10	17:05:01	intranet-1	authlog	N/A	Jan 10 17:05:01 intranet-1 CRON[4507]: pam_unix(cron:session): sessio
2018-01-10	17:05:01	intranet-1	syslog	N/A	Jan 10 17:05:01 intranet-1 CRON[4508]: (root) CMD (command -v debia
2018-01-10	16:55:01	intranet-0	authlog	N/A	Jan 10 16:55:01 intranet-0 CRON[28513]: pam_unix(cron:session): sessi
2018-01-10	16:55:01	intranet-0	authlog	N/A	Jan 10 16:55:01 intranet-0 CRON[28513]: pam_unix(cron:session): sessi
2018-01-10	16:55:01	intranet-0	syslog	N/A	Jan 10 16:55:01 intranet-0 CRON[28514]: (root) CMD (command -v debi
2018-01-10	16:55:01	intranet-2	authlog	N/A	Jan 10 16:55:01 intranet-2 CRON[19870]: pam_unix(cron:session): sessi
2018-01-10	16:55:01	intranet-2	authlog	N/A	Jan 10 16:55:01 intranet-2 CRON[19870]: pam_unix(cron:session): sessi
2018-01-10	16:55:01	intranet-2	syslog	N/A	Jan 10 16:55:01 intranet-2 CRON[19871]: (root) CMD (command -v debi

The following properties are included for each log entry:

- **Date.** The date of the log entry.
- **Time.** The time of the log entry.
- **Host.** The server that originated the log entry.
- **Source.** The log type for the log entry.
- **Severity.** The severity of the log entry.
- **Message.** The log message.

Starting the Logs View

1. Start the **vADC Analytics** application, see [“Accessing the vADC Analytics Application” on page 299](#).
2. Click **Logs** to access the logs.

FIGURE 417 Accessing Logs



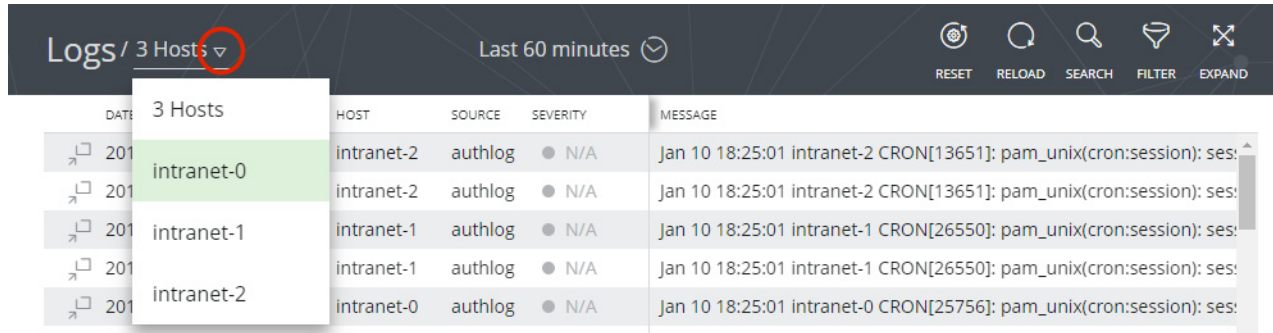
The Logs View appears.

Controlling the Logs View

You can control the display of logs in the following ways:

- You can select a specific originating host for log entries by selecting it from the **Log Filter**.


FIGURE 418 Analytics: Log Filters



To reset the **Log Filter**, select the top listed item. In this example, after selecting the *Intranet-0* host, you can then select *3 Hosts* to revert to using all hosts.


- You can select a time period for displayed logs using the **Time Selector**. This operates in the same way as the **Time Selector** for graph types, see [“Choosing a Time Period” on page 303](#).

FIGURE 419 Analytics: Log Time Period Selector




- You can reset the **Log Filter** at any time by clicking the **Reset** button:

FIGURE 420 Analytics: Resetting Log Filters




- You can refresh retrieved logs by clicking the **Reload** button. For example, to refresh the log data for the *Last 60 minutes*:

FIGURE 421 Analytics: Reloading Logs




- You can search through log entries by clicking the **Search** button. See [“Searching in Displayed Logs” on page 366](#) for full details of this process.

FIGURE 422 Analytics: Searching Logs




- You can configure an extended set of filters in addition to the **Component Filter** by clicking the **Filter** button. This operates in the same way as the **Extended Filter** for graph types, see [“Working with the Extended Filter” on page 316](#).

FIGURE 423 Analytics: Accessing the Extended Filter



- You can maximize the use of space within the browser by clicking the **Expand** toggle.

FIGURE 424 Analytics: Maximizing the Data Area for Logs




Searching in Displayed Logs

You can search through the current displayed log entries using a text string.


- To start a search, click **Search**.

FIGURE 425 Analytics: Searching Logs



The search text box appears.

FIGURE 426 Analytics: The Search Text Box



- Specify a search string. Searches are case-insensitive, and the following special characters are supported:
 - ***: A star matches zero or more characters, excluding whitespace unless the term is enclosed in double quotes. For example, use `*.*.*` to search for log entries that contain an IPv4 address.
 - "**: Use double quotes to enclose one or more spaces within a search term. For example, to search for the phrase *session closed* rather than log entries that contain the words *session* and *closed*, specify `"session closed"`.
 - : A minus sign, used at the start of a search term (outside the double quotes if used), excludes all lines that contain the term. For example:
 - To search for log entries that do not contain *cron*, specify `-cron`.


- To search for log entries that contain *session* but which do not contain *closed*, specify *session -closed*".
- To search for log entries that do not contain the phrase *session closed*, specify *-"session closed"*.
- ****: A backslash can be used to escape all special characters, including ***, *"*, *-*, and itself.

For example, to search for *-logind*, specify *\-logind*

Note: You can view this information by clicking the information button next to the search text box.

3. Press *Enter* or click the lens to search. For example:

FIGURE 427 Analytics: Confirming a Search String



The space-separated terms are then OR-ed together, except for negated terms which are AND-ed with the result of the non-negated terms. For example, to search for the word *closed* in a line that does not also contain the word *session*, specify *session -closed*.


After searching, the number of matching log entries is displayed and matching phrases are highlighted.

FIGURE 428 Analytics: Confirming a Search String

SEARCH					
session -closed		21 found	PREVIOUS	NEXT	
DATE	TIME	HOST	SOURCE	SEVERITY	MESSAGE
2018-01-10	18:25:01	intranet-2	authlog	N/A	Jan 10 18:25:01 intranet-2 CRON[13651]: pam_unix(cron:session): ses:
2018-01-10	18:25:01	intranet-2	authlog	N/A	Jan 10 18:25:01 intranet-2 CRON[13651]: pam_unix(cron: session): se
2018-01-10	18:25:01	intranet-1	authlog	N/A	Jan 10 18:25:01 intranet-1 CRON[26550]: pam_unix(cron:session): ses:
2018-01-10	18:25:01	intranet-1	authlog	N/A	Jan 10 18:25:01 intranet-1 CRON[26550]: pam_unix(cron: session): se
2018-01-10	18:25:01	intranet-0	authlog	N/A	Jan 10 18:25:01 intranet-0 CRON[25756]: pam_unix(cron:session): ses:
2018-01-10	18:25:01	intranet-0	authlog	N/A	Jan 10 18:25:01 intranet-0 CRON[25756]: pam_unix(cron: session): se
2018-01-10	18:25:01	intranet-0	syslog	N/A	Jan 10 18:25:01 intranet-0 CRON[25757]: (root) CMD (command -v del

4. Click **Next** and **Previous** to navigate the located results.
5. (Optional) Click the **Clear** control to reset the search string. For example:

FIGURE 429 Analytics: Clearing a Search String



Working with High Availability

- **Overview: High Availability on Services Director** 369
- **Creating a High Availability Pair in the Services Director VA** 371
- **Viewing High Availability Status** 372
- **Taking a Backup of Your Services Director** 374
- **Responding to Reported Health Issues** 375
- **Swapping the Roles of the HA Nodes** 377
- **Ejecting a Node from an HA Pair** 381
- **Recovering from a Failed Active Node** 385
- **Recovering from a Split Brain Scenario** 389
- **Converting an Ejected Node into a Standalone Active Node** 395
- **Converting an Upgraded Node into a Standalone Active Node** 398

Overview: High Availability on Services Director

High Availability (HA) is a Services Director configuration.

An HA configuration enables two Services Director nodes to operate as a synchronized *HA pair*, with an *Active* Services Director being backed up by a *Standby* Services Director.


Note: The Services Director HA pair and its Service Endpoint Address can be in a private network behind a NAT device.

Each node in the HA pair maintains a database that stores management metadata for various components, including all registered/deployed Virtual Traffic Managers (vTMs) in the network.

The metadata is synchronized from the *Active* node to the *Standby* node.

The HA pair has a Service Endpoint Address (SEA), which points to whichever of the Services Directors is currently the *Active* node. This enables users to always access the Services Director VA using the same hostname/IP address at all times.

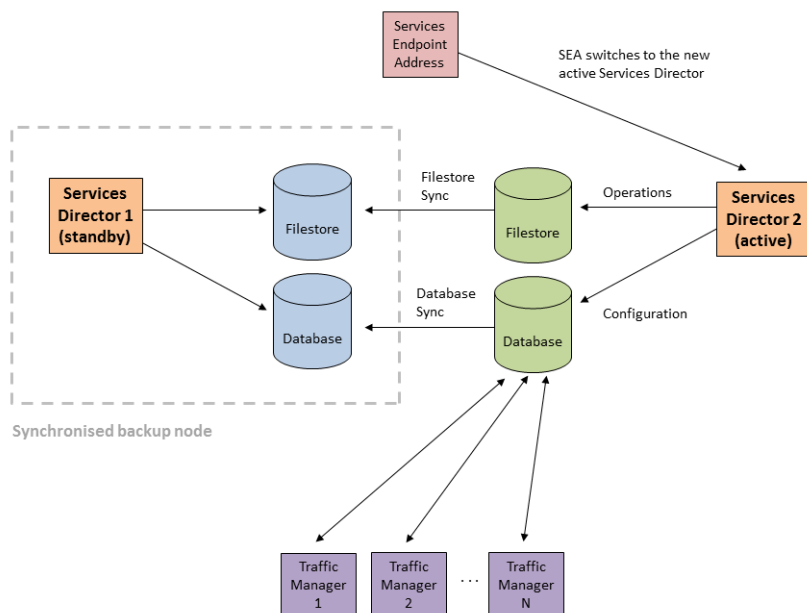
FIGURE 430 High Availability Overview



In the event of failure of the *Active* node, the *Standby* node contains a synchronized copy of the current configuration for the Services Director, and can take over as the *Active* node. The former *Active* node becomes the *Standby* node, and the direction of all synchronization reverses.

The switching process, called *failover*, is triggered manually by the administrator.

FIGURE 431 High Availability Overview: After Failover




Creating a High Availability Pair in the Services Director VA

In the Services Director VA, an HA pair is formed by joining a Secondary Services Director to an existing Primary Services Director.

This process happens during the Setup Wizard for a Secondary Services Director. See [“Installing and Configuring a Secondary Services Director” on page 105](#).

FIGURE 432 High Availability: Creating and Joining Services Directors



Once the HA pair is formed, the concepts of Primary and Secondary Services Directors are largely put aside; these represent the *virtual machine* implementations of the Services Directors, each of which can be uniquely identified by an IP address or a DNS hostname.

Note: The Services Director HA pair and its Service Endpoint Address can be in a private network behind a NAT device.

The concepts of Primary and Secondary are less important than the *role* that each Services Director performs in the HA pair. The supported roles are:

- The *Active* role - the Services Director controls the HA pair for:
 - Web Service. That is, it controls use of the REST API and licensing.
 - Database and Database Synchronization. The system configuration is contained in a database on the *Active* node, and synchronizes to the *Standby* node.
 - File System and File System Synchronization. The file system of the *Active* node is synchronized to the *Standby* node.
- The *Standby* role - the Services Director receives system information from the *Active* node:
 - The synchronized database.
 - The synchronized file system.

The *Active* and *Standby* roles can be changed using software operations, without regard for whether each node is operating on the Primary or Secondary Services Director. See [“Swapping the Roles of the HA Nodes” on page 377](#).

The Service Endpoint Address is the management address for the Services Director as a whole, and always points to the *Active* Services Director node.

Viewing High Availability Status

The current HA status for the Services Director HA pair is shown on the **Services > Manage HA** page of the Services Director VA.

FIGURE 433 Manage HA Page

Manage HA

The screenshot displays the 'Manage HA' page with two side-by-side panels. The left panel, titled 'gold-01 10.62.167.199', has a white background and shows the node is 'Active', handling all service requests, with a Service Endpoint Address of 10.62.167.201. Its health status for 'Web service' and 'Database replication' is 'Good'. The right panel, titled 'silver-01 10.62.167.200', has a blue/gray background and shows the node is 'Standby', not handling any service requests. Its health status is also 'Good'. At the bottom of the right panel, there is an 'Eject' button and a warning: 'Ejecting this node will remove it from the current HA pair. As this node is Standby, the Active node will continue to service all requests.'

The HA pair is represented by a pair of panels on the **Manage HA** page. Each panel shows information for either the *Active* or the *Standby* node.

- The node you are logged in to is always presented on the left.

In this example, you are logged in to the gold-01 node.

- The *Active* node is always presented in a white panel.

In this example, gold-01 is the *Active* node.

- The *Standby* node is always presented in a blue/gray panel.

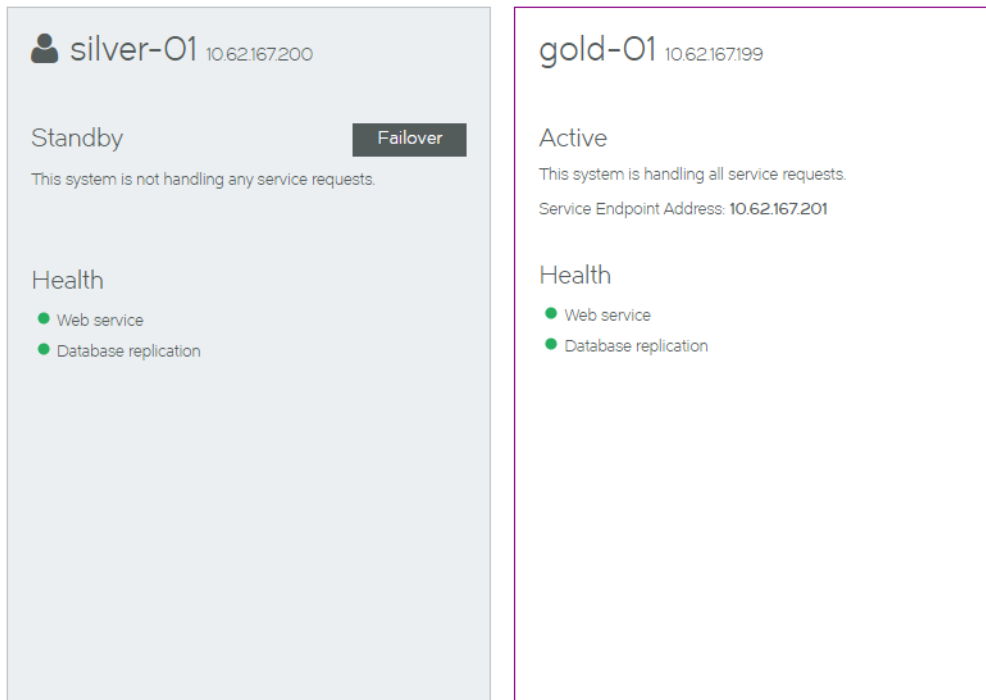
In this example, silver-01 is the *Standby* node.

- Where additional actions are supported, a button is shown.

In this example, the **Eject** button is present on the *Standby* node.

If you are logged in to the *Standby* node, your view will be similar to the following:

FIGURE 434 Manage HA Page: Logged in to Standby Node



Each panel includes health indicators for the node. These indicate the health of:

- Web Services. That is, the REST API services and vTM licensing.
- Database replication.

While an indicator is green, it is healthy.

When one or more of these operations is unhealthy, it is orange. See [“Responding to Reported Health Issues” on page 375](#).

If the Services Director HA pair is in a private network behind a NAT device, the internal Service Endpoint Address and the external IP Address for the HA pair are displayed. For example:

FIGURE 435 Manage HA Page: HA Pair in a Private Network Behind a NAT Device

Manage HA

gold-01 192.168.21129

Active
This system is handling all service requests.

Service Endpoint Address: **192.168.21131**
External IP Address: **10.62.150.31**

Health

- Web service
- Database replication

silver-01 192.168.21130

Standby
This system is not handling any service requests.

Health

- Web service
- Database replication

Eject

Taking a Backup of Your Services Director

When your Services Director system is fully configured, you can preserve its configuration by taking regular scheduled backups. This serves two purposes:

- In the event of a failure of a node's configuration, you can use a backup to recover the configuration.
- In the event of a failure of a Services Director node, you can use a backup to create a new Services Director. This is achieved by using a backup configuration during the Setup Wizard.

See **"Recovering from a Services Director Failure" on page 401** for full details of both scenarios.

Responding to Reported Health Issues

When a node is in an unhealthy state, an orange health indicator is used. For example:

FIGURE 436 Manage HA Page: Unhealthy State

Manage HA

gold-01 10.62.167.199

Active

This system is handling all service requests.

Service Endpoint Address: 10.62.167.201

Health

Web service

Database replication

silver-01 10.62.167.200

Standby

This system is not handling any service requests.

Health

Problems detected.

Web service

Database replication

Diagnose

Eject

Click the **Diagnose** button to understand more about the problem. For example:

FIGURE 437 Manage HA Page: Diagnosing Unhealthy State

Active

This system is handling all service requests.

Service Endpoint Address: 10.62.167.201

Health

Web service

Database replication

Standby

This system is not handling any service requests.

Health

Problems detected.

Web service

Database replication

Diagnose

Eject

Diagnostics


Monitoring of database node has failed.

Close

Several kinds of errors can be reported:

- Some errors are caused by transient issues in your network, and will clear once the network recovers.
If an error does not clear in a few minutes, further investigation may be required.

FIGURE 438 High Availability: Transient Network Issues



- Some errors may require an Administrator to log in to the affected node directly to analyze and fix a reported issue using a reboot, the REST API or the Command-Line User Interface (CLI). Refer to the *Pulse Services Director Advanced User Guide* and the *Secure Services Director Command Reference* for details.
- Some errors are caused by the failure of one of the nodes. To respond to this, you can change the *Active* and *Standby* roles using software operations:
 - The *Standby* node can perform a *failover*. This operation swaps the roles performed by the paired Services Director. Both nodes must be healthy to do this, you must repair the unhealthy node first. Failover is commonly used before performing maintenance on an *Active* node. (see [“Swapping the Roles of the HA Nodes” on page 377](#)).
 - The *Active* node can *eject* an unhealthy *Standby* node in the event of failure. This creates an *Active* standalone Services Director and an unpaired *Standby* Services Director. See [“Ejecting a Node from an HA Pair” on page 381](#).
 - The *Standby* node can perform a *forced failover*. This operation attempts to swap the roles performed by the paired Services Director while the *Active* node is unhealthy. (see [“Recovering from a Failed Active Node” on page 385](#)).
 - An *Active* node can perform a *forced standby* on itself. This operation is used to recover from an exceptional circumstance where both nodes in an HA pair believes itself to be the *Active* node. See [“Recovering from a Split Brain Scenario” on page 389](#).

Swapping the Roles of the HA Nodes

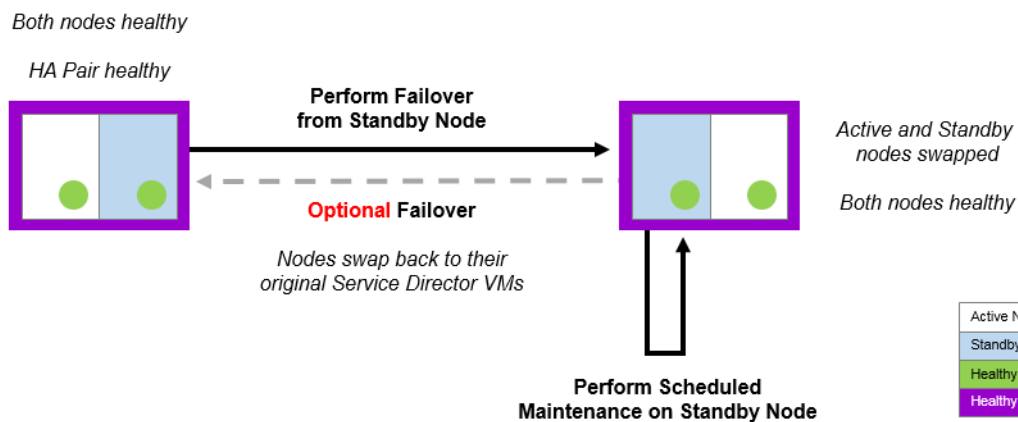
When you swap the roles of the *Active* and *Standby* nodes, the process is called *failover*.

Both nodes must be healthy to perform a failover.

Failover is useful in a number of scenarios:

- Before performing scheduled maintenance on the *Active* node.
- Before performing additional repairs to a recently-repaired *Active* node.
- To enable the current *Active* node to be subsequently ejected.

FIGURE 439 High Availability: Failover



After a failover completes, the Services Endpoint Address points to the new *Active* node.

If either of the nodes is unhealthy, you must repair the unhealthy node first, or use a different operation such as an ejection (see [“Ejecting a Node from an HA Pair” on page 381](#)) or a forced failover (see [“Recovering from a Failed Active Node” on page 385](#)).

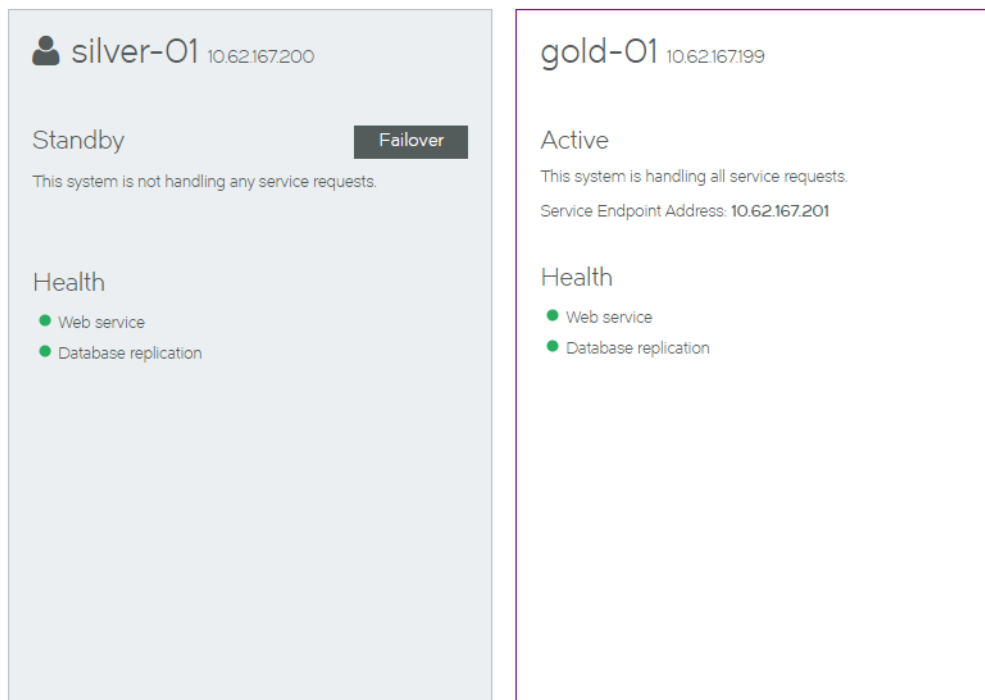
Performing a Failover from the Standby Node

1. Access your *Standby* Services Director VA from a browser, using either the IP address or hostname of your *Standby* node.
2. Log in as the administration user. The **Home** page appears.

- Click the **Services** menu, and then click **High Availability: Manage HA**. The **Manage HA** page appears.

FIGURE 440 Manage HA Page: Before Failover

Manage HA

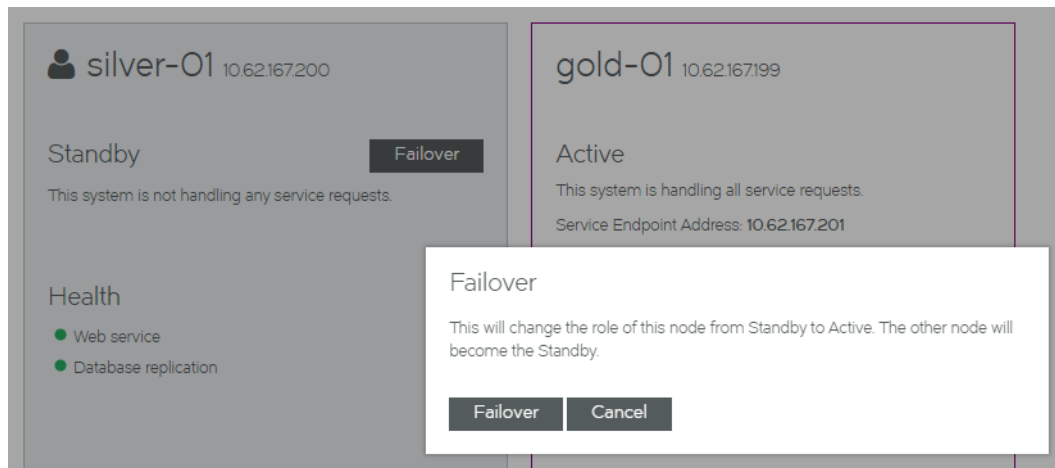


In this example:

- The *Standby* node (silver-01) is displayed on the left in a blue/gray panel.
 - The *Active* node (gold-01) is displayed on the right in a white panel.
 - A **Failover** button is available for the *Standby* node.
- Ensure that all healthy indicators are green.

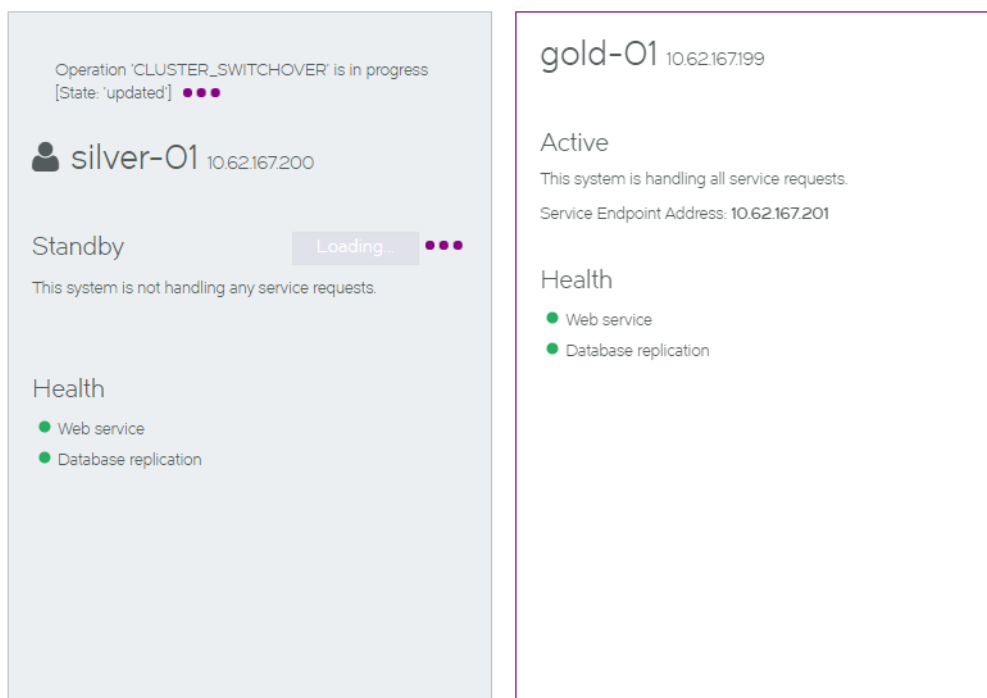
- In the *Standby* panel, click **Failover**. An information panel appears.

FIGURE 441 Manage HA Page: Confirming a Failover



- Click **Failover**. The failover starts.

FIGURE 442 Manage HA Page: Failover In Progress




The failover process reports an error and stops if the *Active* node goes down as the failover is started. A retry of the failover will become a forced failover. See [“Recovering from a Failed Active Node” on page 385](#).

- Wait for the process to complete. The health indicators may become orange during the transition, but these will clear.

After the failover completes, the **Manage HA** page updates:

- the original *Standby* node (silver-01) is now the *Active* node.
- the original *Active* node (gold-01) is now the *Standby* node.
- All health indicators are green.

FIGURE 443 Manage HA Page: Failover Complete



8. (Optional) Perform the following actions


- Perform maintenance on the new *Standby* node.
- Perform another failover to return the Primary Services Director and Secondary Services Director to their original roles.
- Eject the *Standby* node. See [“Ejecting a Node from an HA Pair” on page 381](#).

Ejecting a Node from an HA Pair

A healthy *Active* node can *eject* the other member of an HA pair. This is useful in a number of scenarios:

- Ejecting an unhealthy *Standby* node in the event of failure. This creates a standalone *Active* node and an unpaired unhealthy *Standby* node.

FIGURE 444 High Availability: Ejecting Unhealthy Standby Node




Once the *Standby* node is repaired, it can be joined to any standalone node to form an HA pair.

- Ejecting an unhealthy node after a forced failover operation fails.

In this instance, both nodes are *Active*, but one is unhealthy. The unhealthy *Active* node can be ejected from the healthy *Active* node.

FIGURE 445 High Availability: Ejecting After Forced Failover Fails



- You can also eject a healthy *Standby* node if required. This results in a healthy standalone node and a healthy unpaired node.

Ejecting a Standby Node from the Active Node

1. Access your *Active* Services Director VA from a browser, using the Service Endpoint Address of your Services Director.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **High Availability: Manage HA**. The **Manage HA** page appears.

FIGURE 446 Manage HA Page: Before Ejection

Manage HA

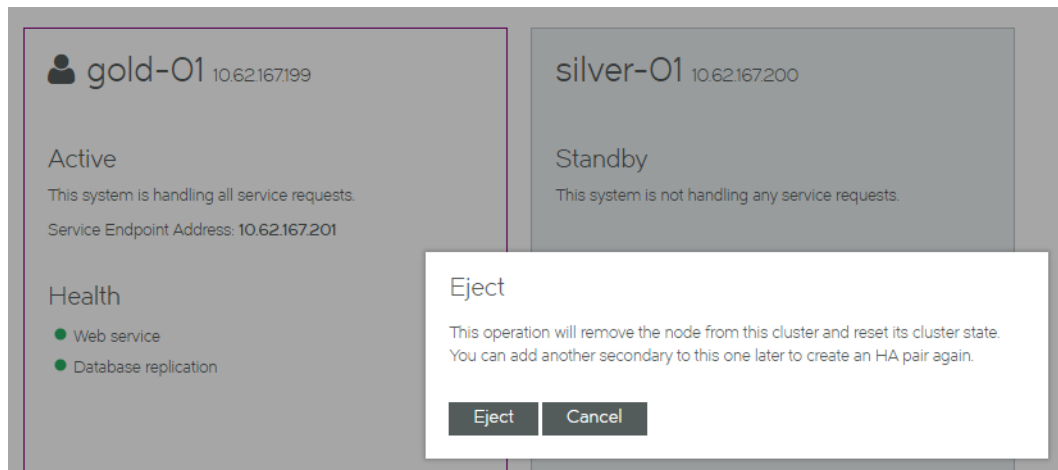
The screenshot displays the 'Manage HA' interface with two panels. The left panel, titled 'gold-01 10.62.167.199', is white and labeled 'Active'. It states 'This system is handling all service requests.' and 'Service Endpoint Address: 10.62.167.201'. Under 'Health', both 'Web service' and 'Database replication' are indicated by green dots. The right panel, titled 'silver-01 10.62.167.200', is blue/gray and labeled 'Standby'. It states 'This system is not handling any service requests.' and also shows green dots for 'Web service' and 'Database replication'. At the bottom of the right panel, a text box explains: 'Ejecting this node will remove it from the current HA pair. As this node is Standby, the Active node will continue to service all requests.' Below this is a dark gray 'Eject' button.

In this example:

- The *Active* node (gold-01) is displayed on the left in a white panel.
 - The *Standby* node (silver-01) is displayed on the right in a blue/gray panel.
 - An **Eject** button is available for the *Standby* node.
4. Ensure that all healthy indicators are green.

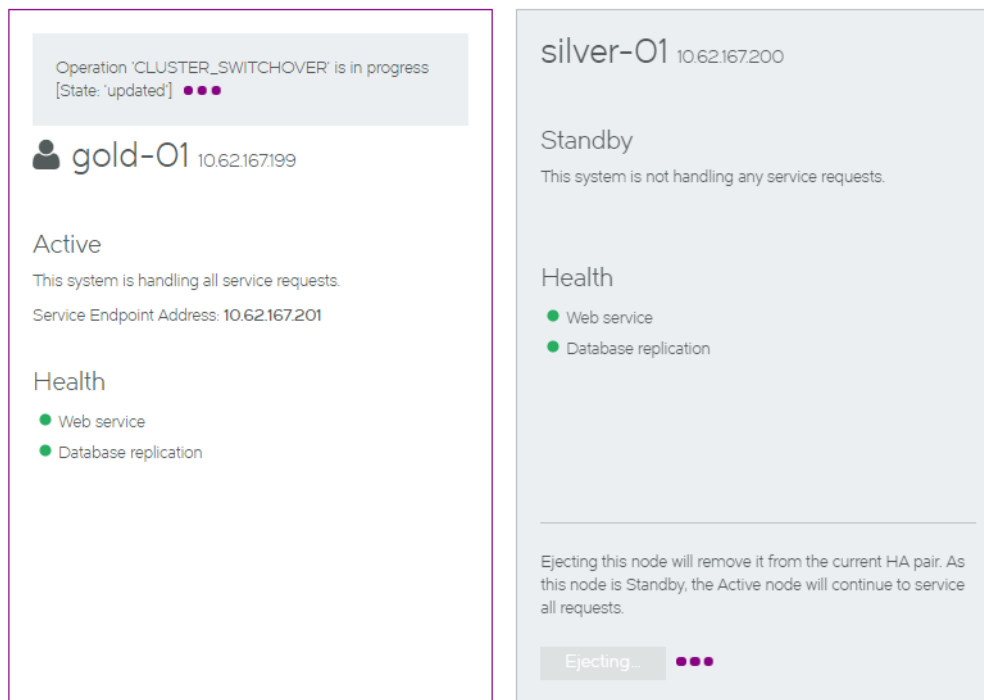
- In the *Standby* panel, click **Eject**. An information panel appears.

FIGURE 447 Manage HA Page: Confirming an Ejection



- Click **Eject**. The ejection starts, and reports progress.

FIGURE 448 Manage HA Page: Ejection In Progress



- Wait for the process to complete. The health indicators may become orange during the transition, but these will clear.

After the ejection completes, the **Manage HA** page updates:

- The original *Active* node (gold-01) remains in place as a standalone node.
- No *Standby* node is configured.

The original *Standby* node still exists, but it is now an unpaired Services Director node.

- All health indicators are green.

FIGURE 449 Manage HA Page: Ejection Complete

gold-O1 10.62.167.199

Active

This system is handling all service requests.

Service Endpoint Address: 10.62.167.201

Health

- Web service
- Database replication

⚠ No secondary configured

Your HA setup has only one node configured. To reduce the chances of service outages it is strongly suggested that you have a pair of nodes in your HA setup. You can do this by joining another node as a secondary to this one.

8. (Optional) Confirm the state of the original *Standby* node. To do this, start its Services Director VA using its IP address or hostname and access its **Manage HA** page.

FIGURE 450 Manage HA Page: Post-Ejection State of Original Standby Node

Manage HA

High availability not configured

To reduce the chances of service outages it is strongly suggested that you enable HA by assigning this node as a HA primary. A primary Services Director can run standalone or paired with the Secondary. When paired with the secondary, the primary will act in an Active role and the secondary will act as a Standby.

Create Primary

From this screen, you can convert this ejected *Standby* node into a standalone *Active* node, see [“Converting an Ejected Node into a Standalone Active Node” on page 395.](#)


Recovering from a Failed Active Node

If your *Active* node becomes unhealthy, it must be repaired.

Maintenance is typically performed on a *Standby* node. However, you cannot perform a failover to swap the *Active* and *Standby* nodes, because a failover requires both nodes to be healthy.

To resolve a failed *Active* node, you must attempt a *forced failover* from the healthy *Standby* node.

FIGURE 451 High Availability: Attempting a Forced Failover



If the forced failover succeeds:

- The healthy *Standby* node becomes the healthy *Active* node.
- The unhealthy *Active* node becomes a *Standby* node.
- You can then perform maintenance on the *Standby* node. Alternatively, you can eject the unhealthy *Standby* node if required (see [“Ejecting a Node from an HA Pair” on page 381](#)).
- The Services Endpoint Address points to the new *Active* node.

If the forced failover fails:

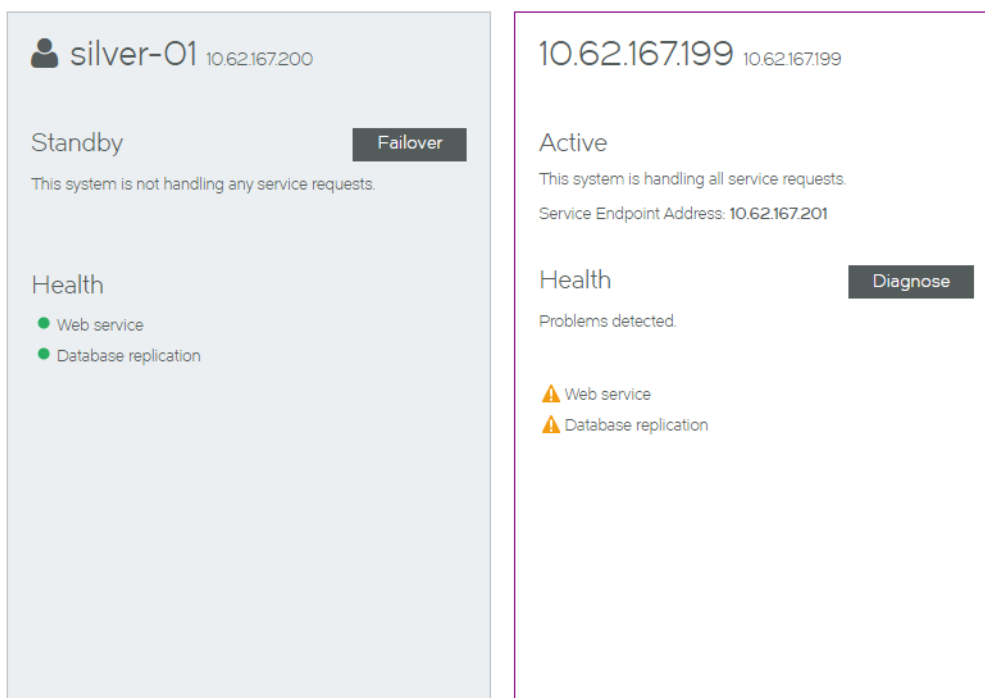
- The healthy *Standby* node becomes a healthy *Active* node.
- The unhealthy *Active* node may remain as an *Active* node. To resolve this you can:
 - Eject the unhealthy *Active* node from the healthy *Active* node (see [“Ejecting a Node from an HA Pair” on page 381](#)).
 - Repair the unhealthy *Active* node. In this case, a "split brain" scenario develops (see [“Recovering from a Split Brain Scenario” on page 389](#)).

To Perform a Forced Failover from the Standby Node

1. Access your *Active* Services Director VA from a browser, using either the IP address or hostname of the healthy *Active* node.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **High Availability: Manage HA**. The **Manage HA** page appears.

FIGURE 452 Manage HA Page: Before Forced Failover

Manage HA



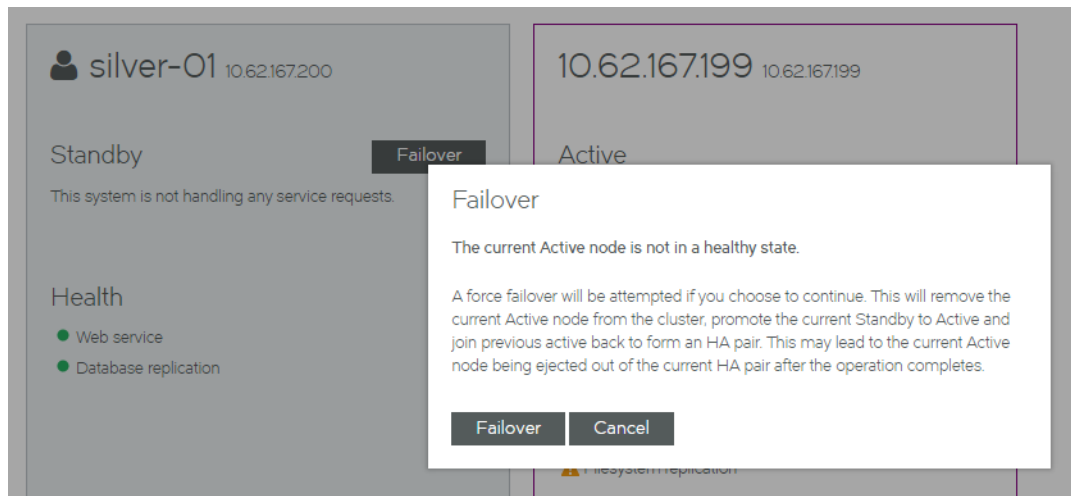
In this example:

- The *Standby* node (silver-01) is healthy.
- The *Active* node (gold-01, identified as 10.62.167.199) is unhealthy.
- The **Failover** button is available for the *Standby* node.

4. Click the **Failover** button.

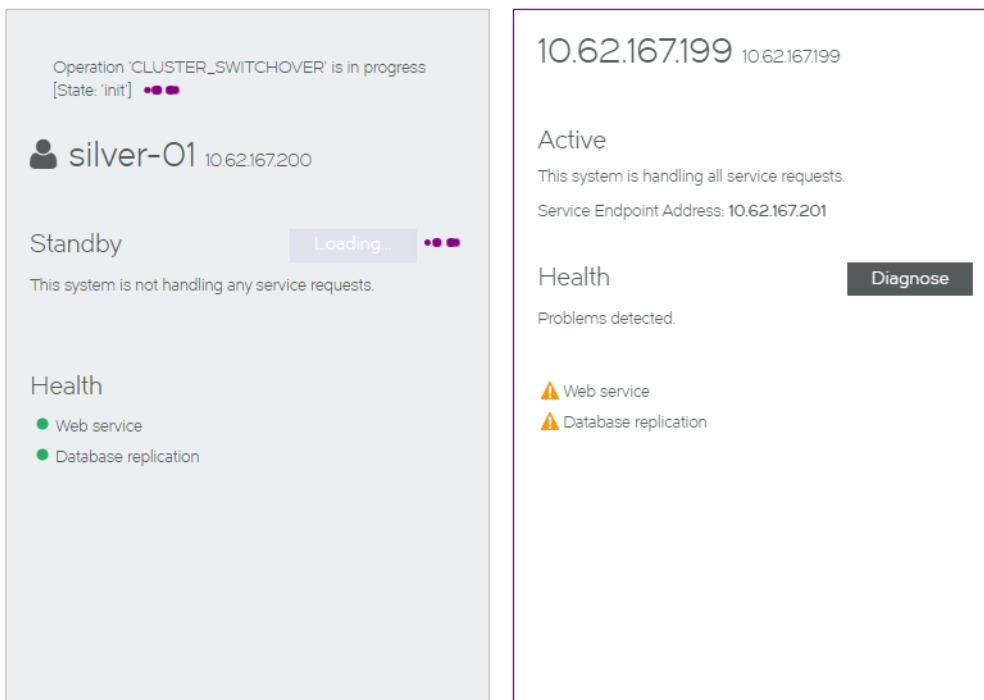
A warning is displayed. This indicates that a forced failover is required, as the *Active* node is not in a healthy state.

FIGURE 453 Manage HA Page: Confirming Forced Failover



- Click **Failover** to confirm the forced failover. The process starts, and displays progress.

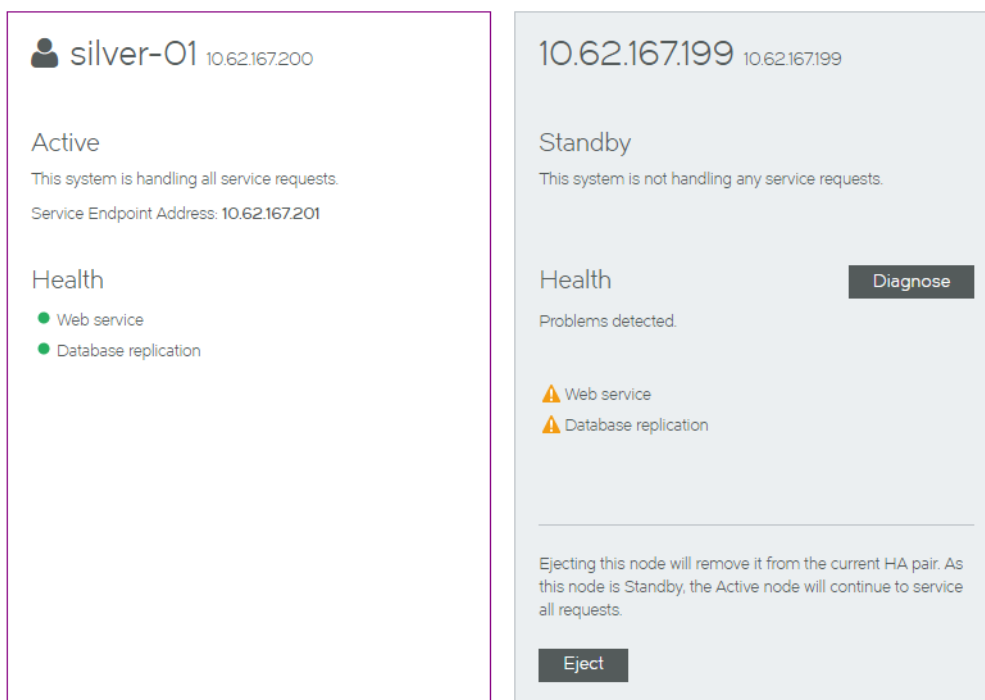
FIGURE 454 Manage HA Page: Forced Failover In Progress



- Wait for the process to complete.

After the ejection completes, the **Manage HA** page updates.

FIGURE 455 Manage HA Page: Forced Failover Complete



It may be difficult to assess the success of this operation from the new *Active* node.

- To assess the success/failure of the forced failover, start the Services Director VA for the unhealthy *Standby* node and access its **Manage HA** page.

If the process has completed successfully:

- The unhealthy *Standby* node is shown on the left
- The healthy *Active* node is shown on the right.

If the process has completed unsuccessfully:

- The unhealthy *Standby* node is shown on the left
- A "split brain" scenario is reported. See [“Recovering from a Split Brain Scenario” on page 389](#) for details.

Recovering from a Split Brain Scenario

The *"split brain"* scenario is an exceptional circumstance where two healthy nodes in an HA pair both believe themselves to be the *Active* node, and that the other node is the *Standby*.


This scenario represents an unhealthy HA pair, and must be resolved.

Understanding How the Split Brain Scenario Arises

The "split brain" scenario can occur after a failed *forced failover* operation. Specifically:

1. The healthy *Standby* node becomes an *Active* node.
2. The unhealthy *Active* node fails to become the *Standby* node.
3. The unhealthy *Active* node is repaired. Both nodes are now healthy and *Active*, and each also believes the other node in the HA pair to be the *Standby* node. This is the "split brain" scenario.

FIGURE 456 High Availability: How "Split Brain" Scenario Occurs



See ["Recovering from a Failed Active Node" on page 385](#) for details of the Forced Failover operation.

Viewing the Split Brain Scenario

A notification of a "split brain" scenario is included in the **Manage HA** page. It is shown in the panel for the *Active* node, along with a **Force Standby** button.

FIGURE 457 High Availability: Notification of "Split Brain" Scenario

Manage HA


The screenshot displays the 'Manage HA' interface with two panels. The left panel, for 'silver-01' (10.62.167.200), is highlighted with a purple border and shows an 'Active' status. It includes a 'Force Standby' button and a notification box stating: 'There seems to be two active nodes in the HA pair. This could happen if the remote node had failed-over to take an 'Active' role while this node was offline. You can make this node a 'Standby' by clicking the button below.' The right panel, for 'gold-01' (10.62.167.199), shows a 'Standby' status and an 'Eject' button. Both panels list 'Web service' and 'Database replication' as healthy components.

Node	IP Address	Status	Health	Action
silver-01	10.62.167.200	Active	Web service, Database replication	Force Standby
gold-01	10.62.167.199	Standby	Web service, Database replication	Eject

Resolving a Split Brain Scenario

To resolve the "split brain" scenario, perform a forced standby operation from the repaired *Active* node. This forces the repaired *Active* node to become the *Standby* node in the HA pair.

FIGURE 458 High Availability: Resolving a "Split Brain" Scenario



1. Access the Services Director VA for the repaired *Active* node from a browser, using either the IP address or hostname of your repaired *Active* node.

Do not access the Services Director VA using the Service Endpoint Address.

2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **High Availability: Manage HA**. The **Manage HA** page appears.

A notification of a split brain is included in the panel for the *Active* node, along with a **Force Standby** button.

FIGURE 459 Manage HA Page: Before Forced Standby

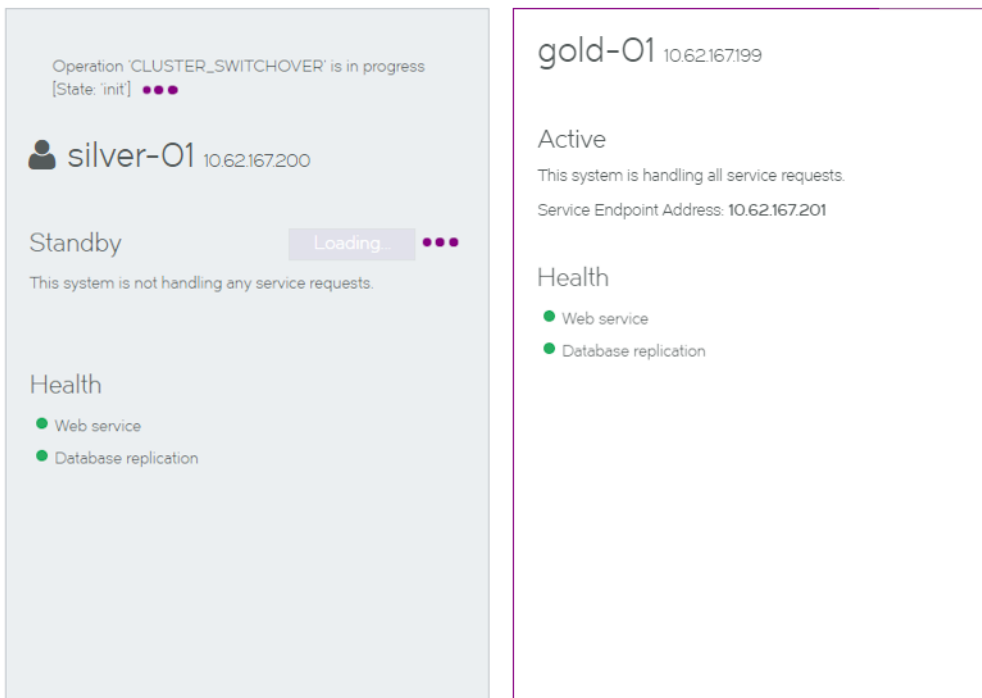
Manage HA

The screenshot displays the 'Manage HA' interface with two panels. The left panel, titled 'silver-01 10.62.167.200', shows the node is 'Active' and handling all service requests. It lists health status for 'Web service' and 'Database replication' as green. A message box states: 'There seems to be two active nodes in the HA pair. This could happen if the remote node had failed-over to take an 'Active' role while this node was offline. You can make this node a 'Standby' by clicking the button below.' Below this is a 'Force Standby' button. The right panel, titled 'gold-01 10.62.167.199', shows the node is 'Standby' and not handling any service requests. It also lists health status for 'Web service' and 'Database replication' as green. Below this is an 'Eject' button.

4. Click **Force Standby**. The forced standby starts, and progress is reported. During this process:
 - The repaired *Active* (in this case, silver-01) becomes the *Standby* node.

- The other *Active* node becomes correctly identified and colored.

FIGURE 460 Manage HA Page: Forced Standby In Progress



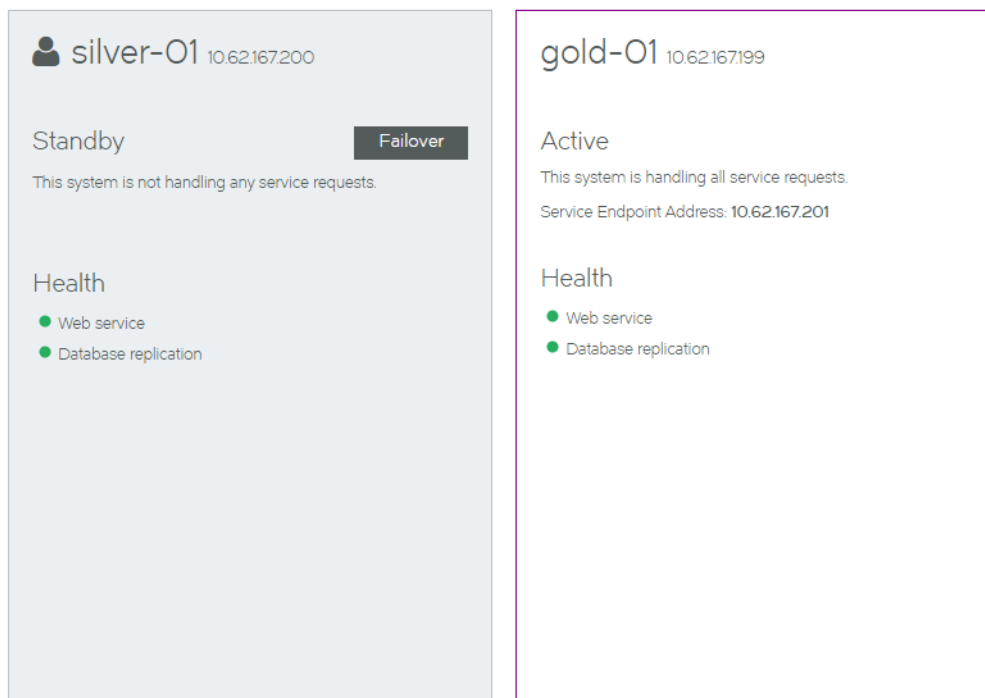
5. Wait for the process to complete. The health indicators may become orange during the transition, but these will clear.

After the forced standby completes, the **Manage HA** page updates:

- The new *Standby* node (silver-01) is on the left.
- The *Active* node (gold-01) is on the right.


- All health indicators are green.

FIGURE 461 Manage HA Page: Force Standby Complete (Standby Node)



6. (Optional) Log out of the *Standby* node and start the Services Director VA for the *Active* node. The **Manage HA** page for this node confirms the correct configuration of nodes following this operation.

FIGURE 462 Manage HA Page: Force Standby Complete (Active Node)



Converting an Ejected Node into a Standalone Active Node

After you have ejected a node, it becomes an unpaired Services Director node. This node contains no configuration or licenses.

You can convert this unpaired node to be a Primary Services Director node if required.

To do this, you must choose how you want the IP address of the node to be used:

- The current management IP address of the node can be used as its new Service Endpoint Address. This requires you to enter a new management IP address for the node.
- The current management IP address of the node will be retained. This requires you to enter a new Service Endpoint Address for the node.

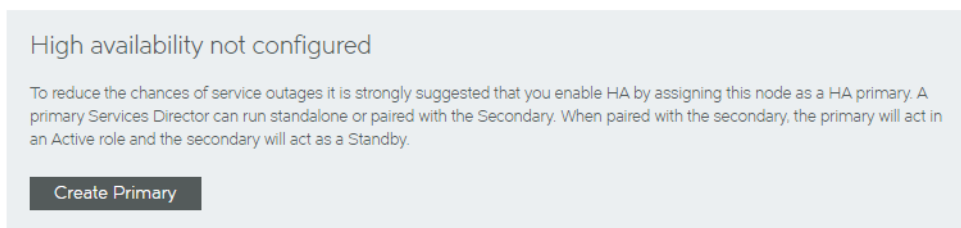
If the Service Endpoint Address is in a private network behind a NAT device, you must also specify the external IP address for the Service Endpoint Address.

1. Access the Services Director VA for the *Standby* node from a browser, using either the IP address or hostname of the *Standby* node.
2. Log in as the administration user. The **Home** page appears.
3. Click the **Services** menu, and then click **High Availability: Manage HA**. The **Manage HA** page appears.

This page confirms the unpaired state of this Services Director node.

FIGURE 463 Manage HA Page: Creating a Standalone Node

Manage HA



4. Click **Create Primary**.

The **Manage HA** page updates to collect the required information.

FIGURE 464 Manage HA Page: Establishing a New Service Point Address

Manage HA

Create a primary HA node

Choose a **Service Endpoint Address**. This address will be used to ensure high-availability as in the event of a failover the secondary services director will be available via the same IP as the primary was accessible from. The service endpoint address must be in the same subnet as the IP on the primary interface.

☒ **Use the IP of the primary interface**

Since the service endpoint address can change from one node to another during a failover, you would need a persistent IP on the primary interface for this node. Please supply a new IP address for the primary interface and a new hostname for this node (hostname that the new IP corresponds to).

NOTE Changing the hostname and IP will take effect immediately and will require navigating back to this page with the new hostname.

Hostname:

Primary interface IP:

☐ **Enter a new service endpoint address**

Service endpoint address:

Service Endpoint Address Type

☒ **The Service Endpoint Address is globally addressable**

☐ **The Service Endpoint Address is behind a NAT device**

External IP Address:

Create

Revert

5. If you want the current management IP address of the node to be used as its new Service Endpoint Address:

- Select **Use the IP of the Primary Interface**.
- Enter a **Hostname** for the new Primary management IP address.
- Enter the new **Primary interface IP** of the node.

This will replace the current management IP address of the node.

6. If you want the current management IP address of the node to be retained:

- Select **Enter a new service endpoint address**.
- Enter a new **Service endpoint address** for the node.

The current management IP address for the node is retained.

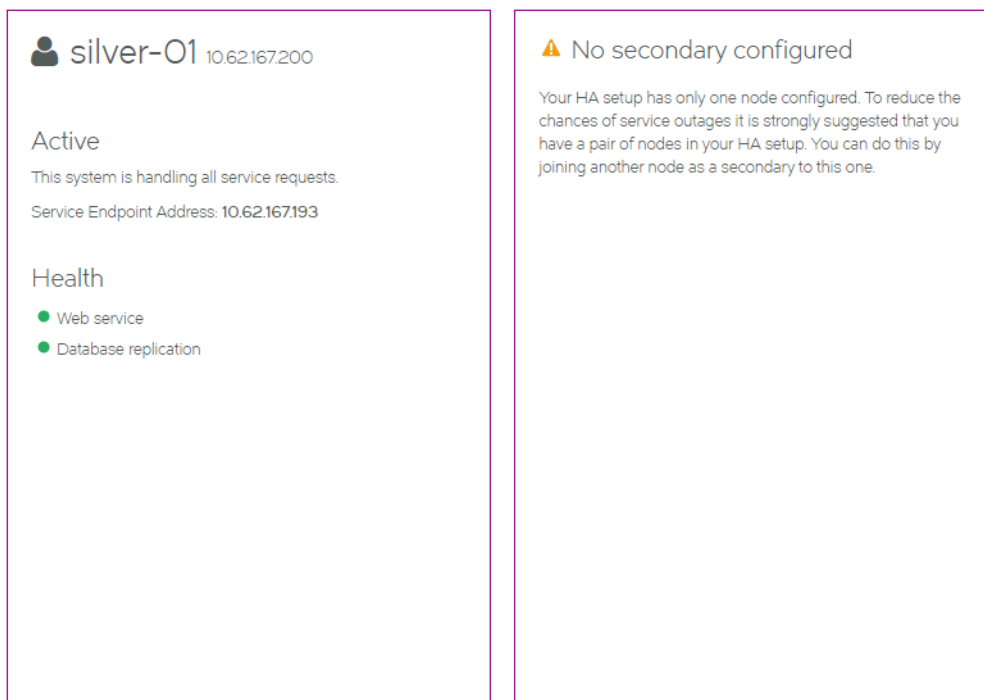
7. If the specified Service Endpoint Address for the Services Director HA pair is globally addressable, select **The Service Endpoint Address is globally addressable**.
8. If the specified Service Endpoint Address for the Services Director HA pair is in a private network behind a NAT device:
 - Select **The Service Endpoint Address is behind a NAT device**.

The available properties update to include an **External IP Address** property.

 - Enter the external NAT address for your Services Director HA pair as the **External IP Address**.
9. Click **Create**. The process starts and reports progress.

When the process completes, the original node is now a standalone Primary Services Director.

FIGURE 465 Manage HA Page: Standalone Node Created



In this example:

- silver-01 retains its original IP address (10.62.167.200)
- silver-01 has a new Service Endpoint Address defined (10.62.167.193).
- silver-01 is now a standalone Primary Services Director.
- silver-01 is not behind a NAT device.

Converting an Upgraded Node into a Standalone Active Node

After you have upgraded your Services Director from an earlier release, it exists as an unpaired Services Director node. This node contains the configuration from the upgraded system.

You can convert this unpaired node to be an Primary Services Director node if required. This enables you to subsequently establish your upgraded node as part of an HA pair.

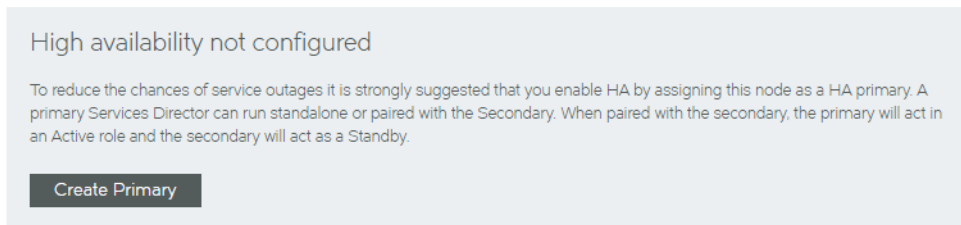
To do this, you will provide the following IP addresses:

- The IP address of your upgraded node becomes the Service Endpoint Address for a standalone Primary Services Director. This ensures that the Legacy FLA licenses that are in use (which must now point to the Service Endpoint Address) will not become invalid during the process.
 - Your upgraded node will then require a new IP address for its management interface.
 - If the Service Endpoint Address is in a private network behind a NAT device, you must also specify the external IP address for the Service Endpoint Address.
1. Access your Services Director VA for the upgraded node from a browser, using either the IP address or hostname of your *Standby* node.
 2. Log in as the administration user. The **Home** page appears.
 3. Click the **Services** menu, and then click **High Availability: Manage HA**. The **Manage HA** page appears.

This page confirms the unpaired state of this Services Director node.

FIGURE 466 Manage HA Page: Creating a Standalone Node

Manage HA



4. Click **Create Primary**.

The **Manage HA** page updates to collect the required information.

FIGURE 467 Manage HA Page: Establishing a New Service Point Address

Manage HA

Create a primary HA node

Choose a **Service Endpoint Address**. This address will be used to ensure high-availability as in the event of a failover the secondary services director will be available via the same IP as the primary was accessible from. The service endpoint address must be in the same subnet as the IP on the primary interface.

☒ **Use the IP of the primary interface**

Since the service endpoint address can change from one node to another during a failover, you would need a persistent IP on the primary interface for this node. Please supply a new IP address for the primary interface and a new hostname for this node (hostname that the new IP corresponds to).

NOTE Changing the hostname and IP will take effect immediately and will require navigating back to this page with the new hostname.

Hostname:

Primary interface IP:

☐ **Enter a new service endpoint address**

Service endpoint address:

Service Endpoint Address Type

☒ **The Service Endpoint Address is globally addressable**

☐ **The Service Endpoint Address is behind a NAT device**

External IP Address:

Create

Revert

5. Select **Use the IP of the Primary Interface**.

6. Enter a **Hostname** for the new Primary management IP address.

This ensures that the current management IP address of your upgraded node becomes its Service Endpoint Address.

7. Enter the new **Primary interface IP** for your upgraded node.

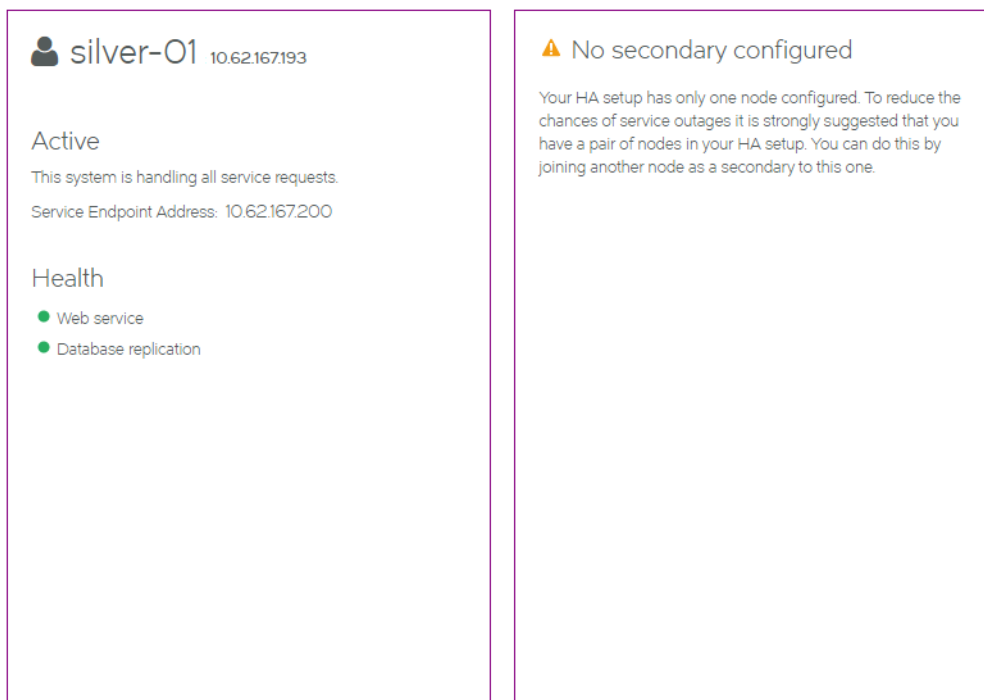
This will replace the current management IP address of your upgraded node.

8. If the Service Endpoint Address for the Services Director HA pair is globally addressable, select **The Service Endpoint Address is globally addressable**.

9. If the Service Endpoint Address for the Services Director HA pair is in a private network behind a NAT device:
 - Select **The Service Endpoint Address is behind a NAT device**.
The available properties update to include an **External IP Address**.
 - Enter the external NAT address for your Services Director HA pair as the **External IP Address**.
10. Click **Create**. The process starts and reports progress.

When the process completes, the original node is now a standalone Primary Services Director.

FIGURE 468 Manage HA Page: Upgraded Node Becomes Standalone Node



In this example:

- silver-01 changes its IP address from 10.62.167.200 to 10.62.167.193.
- silver-01 now has a Service Endpoint Address. This is its original IP address (10.62.167.200).
- silver-01 is now a standalone Primary Services Director. It retains its configuration.
- silver-01 is not behind a NAT device.

When a new Secondary Services Director is created subsequently, it can be joined to silver-01 to form an HA pair. This completes the upgrade process.

Recovering from a Services Director Failure

- **Overview: Recovering from a Services Director Failure** 401
- **Understanding a Backup File** 403
- **Configuring a Scheduled Backup Schedule** 403
- **Restoring a Services Director from a Local Backup** 409
- **Restoring a Services Director from a Remote Backup** 410
- **Restoring a Services Director Using the Setup Wizard** 412
- **Starting and Stopping the Services Director Service** 418


Overview: Recovering from a Services Director Failure

A backup is an encapsulated Services Director configuration. The contents of the backup can be used by the Services Director VA to restore a Services Director configuration.

Backups are made locally according to a backup schedule.

Local backups are copied to a remote server according to a separate schedule.

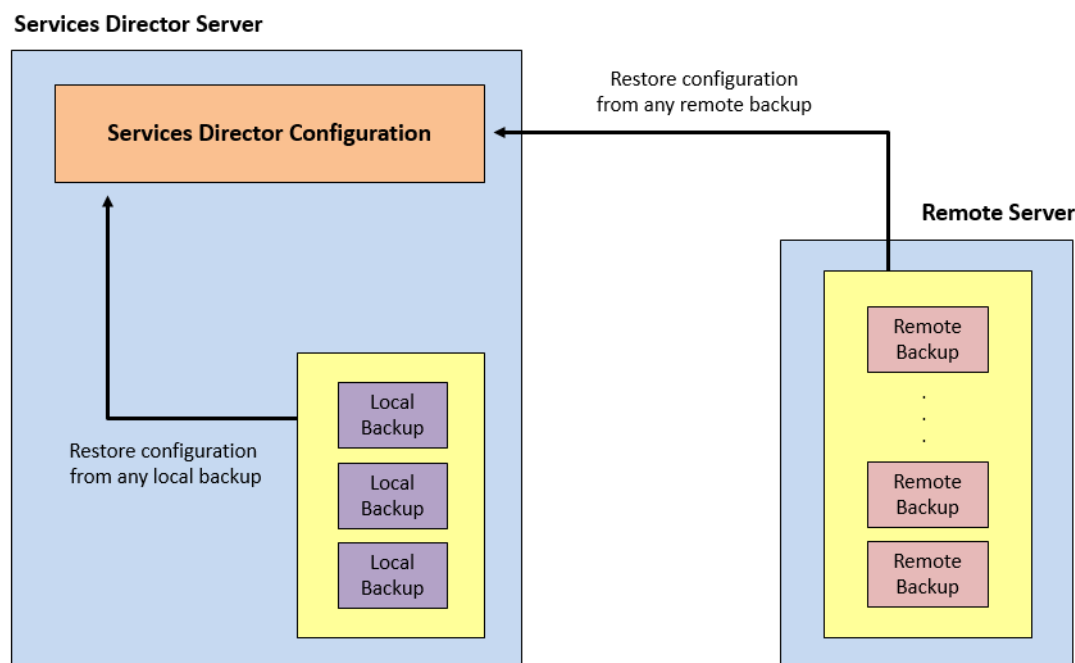
FIGURE 469 Scheduled Backups



Note: Where an HA pair is in use, the backup configuration is created on the *Active* node only. Backups are always restored to an *Active* (or new Primary) node. Standby nodes always take their configuration from the *Active* node.

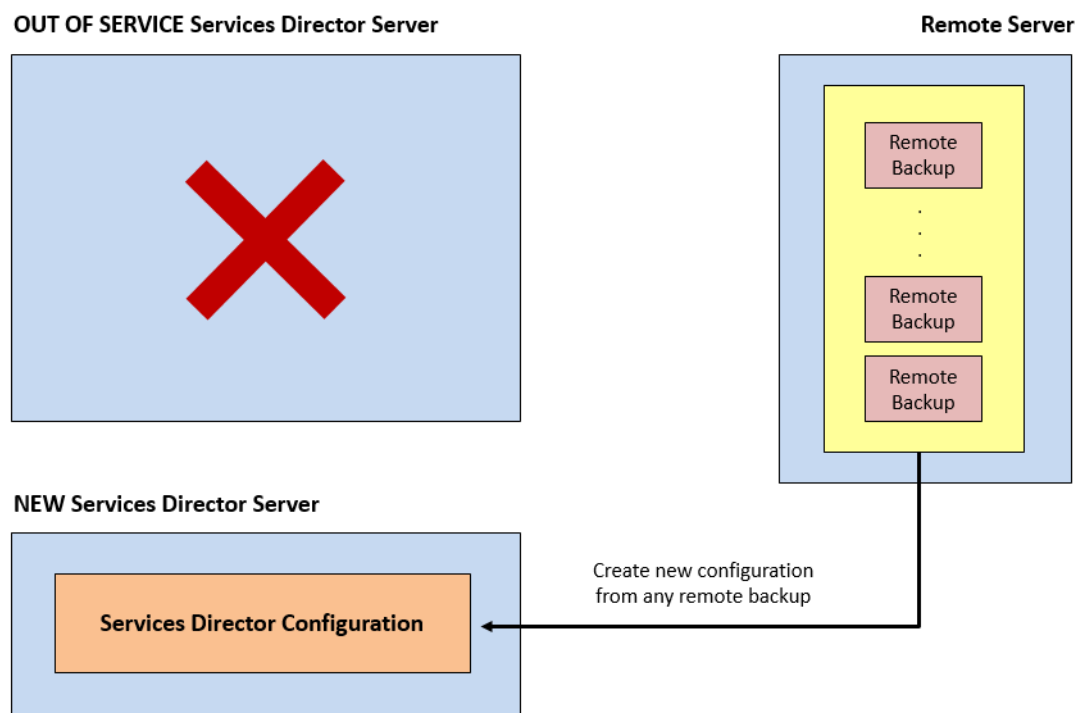
A Services Director VA's configuration can be restored from any backup (either local or remote). You may wish to do this to recover a specific configuration, or to reverse recent changes.

FIGURE 470 Restoring From a Backup



After the failure of a Services Director, a new Services Director VA can be created from the configuration stored in a remote backup.

FIGURE 471 Creating a New Services Director VA From a Backup



Understanding a Backup File

A backup file is a zipped collection of Services Director configuration files. This includes:

- `ssc_build_version.txt`: Version string of VA. For example, 19.1.0-mainline.
- `ssc_certificate.txt`: Certificate and private key used by SD core software, for HTTPS connections.
- `ssc_cfg_backup_mysql_dump_<date>_<time>`: MySQL dump for SD database tables.
- `ssc_cfg_ini.txt`: Configuration snippet of SD core configuration.
- `ssc_fla_license.txt`: List of licenses used by SD. Includes full license strings.
- Universal license and other license files.
- `ssc_mgmt_settings.txt`: Email configuration.
- `user_credentials_node.txt`: Password hash of admin user.

The backup file does *not* include:

- The master password.
- The vTM image files. These must be loaded to both Services Director nodes manually.
- A record of the backup schedule and remote server details.
- SSH keys required for passwordless SSH access.
- Knowledge of HA pairs, hostnames or IP addresses.

Configuring a Scheduled Backup Schedule

The Services Director VA uses a defined backup schedule for a standalone Services Director node or the *Active* node in an HA pair.

Note: Do *not* create a backup schedule from the Standby node in an HA pair. A Standby node always takes its configuration from the *Active* node.

The backup schedule defines:

- The frequency of local backups, and the maximum number of backup files to retain.
- The identity and credentials of a remote file server.

You must set up this remote server before starting the backup configuration process. The server must accept either SCP or FTP connections (or both), and have the required directory structure.

- The frequency of the copy process of local backups to the remote server.

Note: Services Director VA has no influence over the number of backup files stored on the remote server, or the management of these files. This is a user activity outside Services Director VA.

Configuring the Backup Schedule

1. Access your *Active* Services Director VA from a browser, using the Service Endpoint Address of your Services Director.

Do not create a backup schedule from the Standby node in an HA pair. Backups are always created from the *Active* node.

2. Log in as the administration user. The **Home** page appears.
3. Click the **System** menu, and then click **Disaster Recover: Backup and Restore**. The **Backup and Restore** page appears.

FIGURE 472 Backup and Restore Page

Backup and Restore

Configure remote backups

You have not configured any remote backup schedule. It is highly recommended that you configure remote backups to aid with disaster recovery.

Remote backup IP/hostname	<input type="text"/>
Remote backup path	<input type="text"/>
Remote system username	<input type="text"/>
Remote system password	<input type="password"/>
Remote backup protocol	SCP ▼
Take a backup every	<input type="text" value="12"/> Hours ▼
Transfer backups every	<input type="text" value="1"/> Days ▼

Use the setting below to configure the number of the local copies of the backups to be retained.

Retain the last (N) backups locally	<input type="text" value="30"/> ▼
-------------------------------------	-----------------------------------

This example indicates that no backup configuration currently exists.

4. Enter the details for the remote server:

- **Remote backup IP/hostname:** This is the IP address or FQDN of the remote server.
- **Remote backup path:** This identifies a directory on the remote server for the backups. This requires a "full path" directory structure. Relative paths cannot be used.
- **Remote system username:** The user name for the remote server.
- **Remote system password:** The password for the user.
- **Remote backup protocol:** The file transfer protocol for the remote backup server. This is either FTP or SCP. Use SCP for secure encrypted transfers.

5. Define the frequency for the local backup. Under **Take a backup every:**

- Select the units for the backup. This can be *Minutes*, *Hours* (default) or *Days*.
- Enter the number of the selected units.

Minutes can range from 1-59, *Hours* from 1-23 and *Days* from 1-31. The default is 12.

For example: 30 Minutes.

6. Define the frequency for copying local backups to the remote server. This will typically be a longer frequency than the one used for local backups. Under **Transfer backups every:**

- Select the units for the backup. This can be *Minutes*, *Hours* or *Days* (default).
- Enter the number of the selected units.

Minutes can range from 1-59, *Hours* from 1-23 and *Days* from 1-31. The default is 1.

For example: 1 Days.

7. Select the maximum number of local backups as **Retain the last (N) backups locally**. The default is 30. This value must be at least equal to the number of backups between remote copies, else backup files will be lost.

The most recent backup files are retained. Any older files are deleted if this limit is exceeded.

FIGURE 473 Backup and Restore Page: Completed

Backup and Restore

Configure remote backups

You have not configured any remote backup schedule. It is highly recommended that you configure remote backups to aid with disaster recovery.

Remote backup IP/hostname	<input type="text" value="10.62.165.128"/>
Remote backup path	<input type="text" value="/home/sd-backup/sd-gold-silver"/>
Remote system username	<input type="text" value="sd-backup"/>
Remote system password	<input type="password" value="....."/>
Remote backup protocol	<input type="text" value="SCP"/> ▼
Take a backup every	<input type="text" value="2"/> <input type="text" value="Hours"/> ▼
Transfer backups every	<input type="text" value="1"/> <input type="text" value="Days"/> ▼

Use the setting below to configure the number of the local copies of the backups to be retained.

Retain the last (N) backups locally	<input type="text" value="30"/> ▼
-------------------------------------	-----------------------------------

- Click **Apply** to confirm the backup schedule.

An empty test file is sent immediately to the remote server.

The backup configuration, including a status indicator, is included on the **Backup and Restore** page.

FIGURE 474 Backup and Restore Page: Healthy Configuration

Backup and Restore

Backup Service Health

Backup Service Health ●

Backing up locally every 2 hours and copying every day to 10.62.165.128:/home/sd-backup/sd-gold-silver [Edit](#)

- Log in to the remote server and ensure that the backup test file is present. If this is not present, check the details for your remote server on the **Backup and Restore** page. An error message will explain the issue.

The first local backup will be created after the full duration of the local backup frequency. For example, after 2 Hours. The file name has the following general form:

backup_<IP_address>_<datestamp>_<timestamp>.zip

For example:

backup_10.62.167.199_2017-09-13_23-32-01.zip

The first copy of local files to the remote server will occur after the full duration of the remote copy frequency. For example, after 1 Days. Any local backup files that are not present on the remote server are copied over.

Updating the Backup Schedule

- Access your *Active* Services Director VA from a browser, using the Service Endpoint Address of your Services Director.


Do not update a backup schedule from the Standby node in an HA pair. Backups are always updated from the *Active* node.

- Log in as the administration user. The **Home** page appears.
- Click the **System** menu, and then click **Disaster Recover: Backup and Restore**. The **Backup and Restore** page appears. This displays a summary of your current backup schedule, and includes a status indicator.

FIGURE 475 Backup and Restore Page: Schedule Summary

Backup and Restore

Backup Service Health


Backup Service Health 

Backing up locally every 2 hours and copying every day to 10.62.165.128:/home/sd-backup/sd-gold-silver [Edit](#)

- Click **Edit** to display the full details.

FIGURE 476 Backup and Restore Page: Editing the Schedule

Backup Service Health

Backup Service Health 

Backing up locally every 2 hours and copying every day to 10.62.165.128:/home/sd-backup/sd-gold-silver [Hide](#)

Remote backup IP/hostname	<input type="text" value="10.62.165.128"/>
Remote backup path	<input type="text" value="/home/sd-backup/sd-gold-silver"/>
Remote system username	<input type="text" value="sd-backup"/>
Remote system password	<input type="password" value="••••••••"/>
Remote backup protocol	<input type="text" value="SCP"/>
Take a backup every	<input type="text" value="2"/> <input type="text" value="Hours"/>
Transfer backups every	<input type="text" value="1"/> <input type="text" value="Days"/>
Use the setting below to configure the number of the local copies of the backups to be retained.	
Retain the last (N) backups locally	<input type="text" value="30"/>

- Make the required changes to your schedule.

Remote backup path requires a "full path" directory structure. Relative paths cannot be used.

- Click **Apply** to confirm the changes.

The first local backup will be created after the full duration of the local backup frequency. For example, after 20 minutes.

The first copy of local files to the remote server will occur after the full duration of the remote copy frequency. For example, after 1 day.

Restoring a Services Director from a Local Backup

A Services Director VA's configuration can be restored from a local backup. You may wish to do this to recover a specific configuration, or to reverse recent changes.

Note: The backup file does not include any vTM image files that you have imported. However, a list of these images is included in the backup, and this list is displayed the end of the process. These must be loaded to both Services Director nodes manually.

1. Access your *Active* Services Director VA graphical interface from a browser, using the Service Endpoint Address of your Services Director.

Do this from a browser, using the Service Endpoint Address of your Services Director.

Do not restore a configuration from the Standby node in an HA pair. Backups are always restored on the *Active* node.

2. Log in as the administration user. The **Home** page appears.
3. Click the **System** menu, and then click **Disaster Recover: Backup and Restore**. The **Backup and Restore** page appears.


This page contains a summary of the current backup schedule, a backup service health indicator, and provides access to the restore functions.

4. Click the **Restore from a local backup** tab.

FIGURE 477 Backup and Restore Page: Restore from a Local Backup

Backup and Restore

Backup Service Health

Backup Service Health 

Backing up locally every 2 hours and copying every day to 10.62.165.128:/home/sd-backup/sd-gold-silver [Edit](#)

Restore from a backup

Restore from a local backup

Restore from a remote backup

Master Password

☐ Store the password to a file

Date and time of backup

2017-09-14 10:30:02 ▼

Restore

5. Enter the **Master Password** that was in place when the backup was taken.
6. Select the required local backup from the pull-down list.

The file names have the following general form:

```
backup_<IP_address>_<datestamp>_<timestamp>.zip
```

7. (Optional) Select the **Store the password to a file** check box to store the master password internally for future use.
8. Click **Restore** to start the restore process.

Once the process completes, the Services Director will be configured in the same way as the original Services Director, including vTMs in its estate.

Note: When the restore completes, any vTM image files referenced in the backup will not be present on your Services Director. You will need to reload them into the **vTM images** page if this is the case.

Note: The vTM image files must be loaded to both Services Director nodes manually.

Refer to the *Pulse Services Director Advanced User Guide* for full details.

Restoring a Services Director from a Remote Backup

A Services Director VA's configuration can be restored from a remote backup. You may wish to do this to recover a specific configuration, or to reverse recent changes.

The Services Director VA is not able to list available backup files on the remote server. You must know the name of the file you wish to restore from before beginning this process.

Note: The backup file does not include any vTM image files that you have imported. However, a list of these images is included in the backup, and this list is displayed at the end of the process. These must be loaded to both Services Director nodes manually.

1. Access your *Active* Services Director VA from a browser, using the Service Endpoint Address of your Services Director.

Do this from a browser, using the Service Endpoint Address of your Services Director.

Do not restore a configuration from the Standby node in an HA pair. Backups are always restored on the *Active* node.

2. Log in as the administration user. The **Home** page appears.
3. Click the **System** menu, and then click **Disaster Recover: Backup and Restore**. The **Backup and Restore** page appears.

This page contains a summary of the current backup schedule, a backup service health indicator, and provides access to the restore functions.

4. Click the **Restore from a remote backup** tab.

FIGURE 478 Backup and Restore Page: Restore from a Remote Backup

Backup and Restore

Backup Service Health

Backup Service Health ●

Backing up locally every 3 hours and copying every 6 hours to 10.62.166.206./space/sd-backup/sd-backup-test/gold-silver-backups [Edit](#)

Restore from a backup

Restore from a local backup

Restore from a remote backup

Restoring from 10.62.166.206./space/sd-backup/sd-backup-test/gold-silver-backups [Edit](#)

Master Password ☐ Store the password to a file

Remote backup filename

5. Enter the **Master Password** that was current when the remote backup was taken.
6. Enter the name of the remote backup file. The file names have the following general form:

`backup_<IP_address>_<datestamp>_<timestamp>.zip`
For example:

`backup_10.62.167.199_2015-09-09_05-52-02.zip`

7. If you want to change the source of the remote backup:
 - a. Click **Edit**. The dialog expands to show additional fields.
 - b. Enter new details for the remote server:
 - **Remote backup IP/hostname** - this is the IP address or FQDN of the remote server.
 - **Remote backup path** - this identifies a directory on the remote server for the backups. This requires a "full path" directory structure. Relative paths cannot be used.
 - **Remote system username** - the user name for the remote server.
 - **Remote system password** - the password for the user.
 - **Remote backup protocol** - the file transfer protocol for the remote backup server. This is either FTP or SCP. Use SCP for secure encrypted transfers.
 - c. Click **Apply** to confirm the changes.
8. (Optional) Select the **Store the password to a file** check box to store the master password internally for future use.

9. Click **Restore** to start the restore process.

Once the process completes, the Services Director will be configured in the same way as the original Services Director, including vTMs in its estate.

Note: When the restore completes, any vTM image files referenced in the backup will not be present on your Services Director. You will need to reload them from the **vTM images** page if this is the case.

Note: The vTM image files must be loaded to both Services Director nodes manually.

Refer to the *Pulse Services Director Advanced User Guide* for full details.

Restoring a Services Director Using the Setup Wizard

After the failure of a Services Director, you can create a new Primary Services Director VA from a remote backup file. This process uses the Setup Wizard. You can then create a new Secondary Services Director VA and pair it with the recovered Primary Services Director VA.

Note: A new Secondary Services Director VA will receive its configuration from the Primary. You do not need to use a restore process when you create the Secondary.

Note that:

- If your new Services Director VA uses a different Service Endpoint Address than the one used for the original Services Director VA, the FLA Licensing of vTM instances will be disrupted.
- A Service Endpoint Address is still required a standalone Primary Services Director. It must be different from the IP address of the Primary Services Director.
- The Services Director VA is unconfigured at this point, and has no record of the remote server. The required backup file must be downloaded from the remote server to the local machine before beginning the backup.

Note: The backup file does not include any vTM image files that you have imported. However, a list of these images is included in the backup, and this list is displayed at the end of the process. These must be loaded to both Services Director nodes manually.

- You require the master password for the original Services Director VA.

Perform the following process:

1. Create a new virtual machine for the Services Director VA using your chosen platform.
2. Start the VM and make a note of its assigned IP address.
3. Access the Services Director VA in a browser window using its IP address.

The Setup Wizard starts.

4. Work through the Setup Wizard until you reach the **Service Endpoint Address** page.

FIGURE 479 Setup Wizard: Service Endpoint Address Page

Pulse Secure

Service Endpoint Address

Choose a Service Endpoint IP address that will be used by this system. The Service Endpoint IP is used to ensure high-availability as in the event of a failover, the Secondary Services director will be available via the same IP address that the Primary was accessible from.

Service Endpoint IP Address

☒ The Service Endpoint Address is globally addressable
☐ The Service Endpoint Address is behind a NAT device

NOTE

After Setup is complete, you should use the **Service Endpoint Address** to locate this system, not the IP used by the network interface in the Network Configuration step. This is also the IP address you should provide to Pulse Secure in order to generate your FLA license (or if you supply a hostname, a hostname which maps to this IP address).

Previous Next

5. If the Service Endpoint Address for the Services Director HA pair is globally addressable:
 - Select **The Service Endpoint Address is globally addressable**.
 - Enter the **Service Endpoint IP Address** for the Services Director HA pair.
6. If Service Endpoint Address for the Services Director HA pair is in a private network behind a NAT device:
 - Select **The Service Endpoint Address is behind a NAT device**. The available properties update to include an **External IP Address**.
 - Enter the internal NAT Service Endpoint Address for your Services Director HA pair as the **Service Endpoint IP Address**.
 - Enter the external NAT address for your Services Director HA pair as the **External IP Address**.

- Click **Next**. The **Restore from Backup** page appears.

FIGURE 480 Setup Wizard: Restore from Backup Page

Pulse Secure®

Restore from backup

If you have a backup file from a previous installation, you can restore it now. Otherwise you can proceed with a new installation.

☒ This is a new system

☐ Restore from a previous backup

[Choose File](#)

Master Password: ☒ Save the password?

NOTE

Backup files do not include vTM images that may have been in use. If you were using managed vTM instances in your previous installation, you will need to re-upload the vTM image files separately after completing Setup.

For security, it is recommended that the master password is input manually every time the Services Director starts. However, the password could be stored in a file (less secure) for non-interactive start up.

[Previous](#) [Next](#)

- Click **Restore from a previous backup**.
- Click **Choose file** and locate the backup file. This file must already be downloaded from the remote server to a local machine. The file names have the following general form:

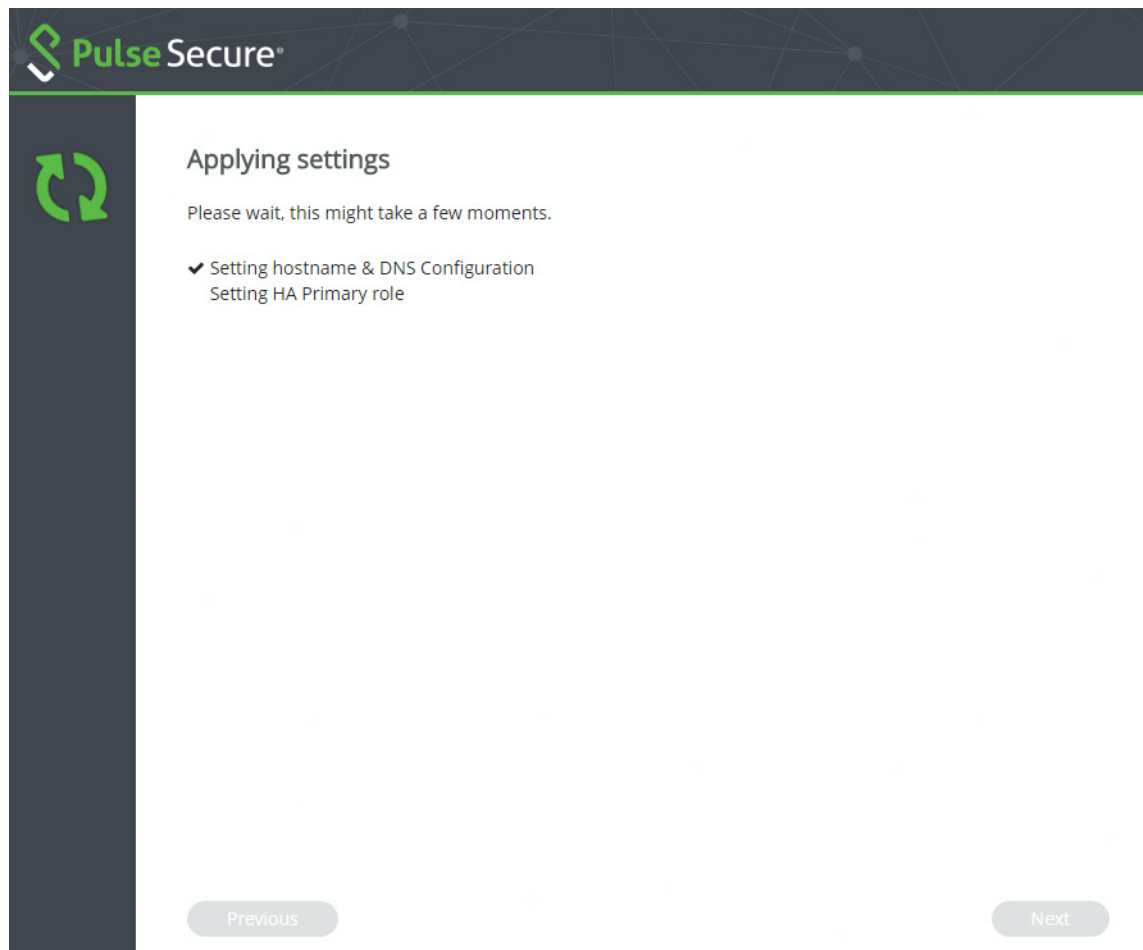
backup_<IP_address>_<datestamp>_<timestamp>.zip

- Enter the Master Password for the Services Director VA that created the backup.

11. Click **Next**. The **Applying Settings** page appears.


This page configures the system based on retrieved configuration information.

FIGURE 481 Setup Wizard: Applying Settings Page



When this is complete, the **Setup Complete** page appears.

FIGURE 482 Setup Wizard: Setup Complete Page



Once the process completes, the Services Director will be configured in the same way as the original Services Director, including vTMs in its estate.

Note: Any vTM image files referenced in the backup will not be present on your Services Director. You will need to reload them from the **vTM images** page if this is the case. These must be loaded to both Services Director nodes manually. Refer to the *Pulse Services Director Advanced User Guide* for full details.

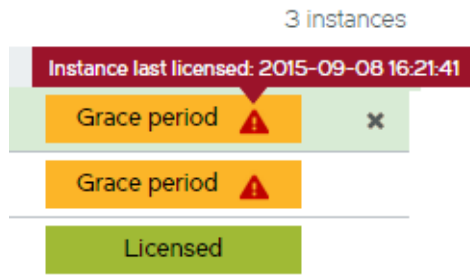
12. Click **Finish**. The **Home** page is displayed.
13. (Optional) Click the **Services** menu, and then click **Services Director: vTM Instances**. The **vTM Instances** page appears.

This page indicates the licensing state of each vTM.

All vTMs that were present in the original configuration should now be present.

If you are using a different Service Endpoint Address to the one used by the FLA Licensing in the backup, the licensing of the vTMs will be disrupted. Each affected vTM will enter a grace period (six weeks). For example:

FIGURE 483 Grace Periods



In this case, generate a FLA license that is keyed to the new Service Endpoint Address. Then, relicense your vTM instances. See [“Relicensing a Virtual Traffic Manager Instance” on page 221](#).

14. Click the **System** menu, and then click **Disaster Recovery: Backup and Restore**. The **Backup and Restore** page appears.

No backup schedule will be present. This information is not saved in the backup.

15. (Optional) Create a new backup schedule. See [“Configuring a Scheduled Backup Schedule” on page 403](#).

The restore process is then complete.

After the restore process is complete for the Primary Services Director VA, you can then create a new Secondary Services Director VA, and join it to the Primary. See [“Preparing to Install the Services Director Virtual Appliance” on page 9](#).

A new Secondary Services Director VA will receive its configuration from the Primary. You do not need to use a restore process when you create the Secondary.

Starting and Stopping the Services Director Service

You can perform a number of master password tasks from the **System** menu.

Restarting the Services Director VA

You can stop, start and restart your Services Director service at any time from the **System > Service Status** page.

- When the system is running, click **Stop** to stop the service.
- When the system is running, click **Restart** to stop and then start the service.
- When the system is not running, click **Start** to start the service.

All changes are immediate.

You are *not* required to enter the master password during this operation. The master password is only required when restarting the Virtual Machine for a Services Director VA. See [“Entering the Master Password After a Virtual Machine Restart” on page 418](#).

Entering the Master Password After a Virtual Machine Restart

You can restart the Virtual Machine (VM) for a Services Director VA at any time.

- If you chose to store the master password internally when you configured the Services Director VA node, you do not need to enter the master password after a VM restart.
- If you did not store the master password internally, you must enter the master password to unlock access to vTMs.

When the Services Director VA is accessed for the first time after a VM restart, the following dialog box appears:

FIGURE 484 Master Password Entry

Services will run in a degraded state until a master password is entered.

Password

☐ I will set the password from the System > Security page later.

There are two scenarios:

- If you know the master password, you will typically enter it immediately. See **“Entering the Master Password Immediately After a Restart” on page 419.**
- If you do not know the master password, but are an administration user, you may want to access the Services Director VA to access functionality that is unrelated to vTMs. For example, to access system logs. You will enter the password at some point afterwards, and regain access to vTM instances. See **“Entering the Master Password Later” on page 419.**

Entering the Master Password Immediately After a Restart

If you know the master password, you will typically enter it immediately.

Note: You may receive an e-mail notification of a raised master_password_fail alarm before you enter the new master password on the Services Director VA.

1. On the master password dialog box, enter the master **Password**.
2. Click **Submit**. This unlocks access to the Services Director VA.
3. To confirm access to vTMs, click the **Services** menu, and then click **Services Director: vTM Instances**. The **vTM Instances** page appears. This page will include all of your vTMs.

Entering the Master Password Later

If you do not know the master password, but are an administration user, you may want to access the Services Director VA to access functionality that is unrelated to vTM instances. Under these circumstances, you can choose to enter the master password at a later point.

Note: If the VM is restarted again, this choice remains in place.

Note: You may receive an e-mail notification of a raised master_password_fail alarm before you enter the new master password on the Services Director VA.

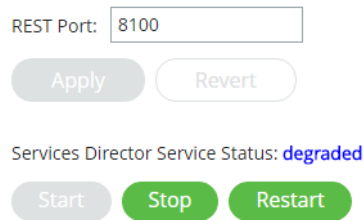
Choosing to Enter the Master Password Later

1. On the master password dialog box, click the **I will set the password from the System Security page later** check box.
2. Click **Submit**.

This unlocks access to the Services Director VA. However, until you enter the master password, the Services Director service status is Degraded. This is indicated on the **System > Service Status** page.

FIGURE 485 Service Status: Degraded Service

Service Status



REST Port:

Services Director Service Status: **degraded**

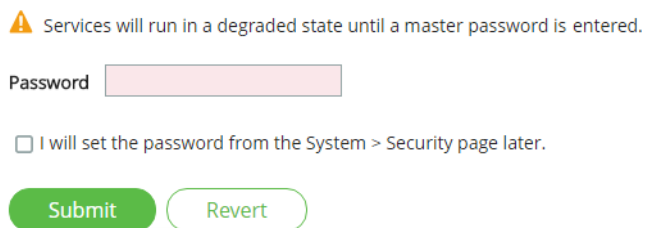
You will have no access to vTMs while in this state.


When you are ready to recover from this Degraded state, you must enter the master password.

Entering the Master Password

1. Click the **System** Menu, then click **Security**. The **Security Settings** page appears.

FIGURE 486 Security Settings Page



 Services will run in a degraded state until a master password is entered.

Password

☐ I will set the password from the System > Security page later.

2. Enter the master password.
3. (Optional) Select the **Store the password to a file** check box to store the master password internally for future use.
4. Click **Submit**.

The **Security Settings** page updates, but no further action is required on this page.

5. Click the **System** menu, then click **Service Status**. The **Service Status** page appears, which enables you to confirm that the Degraded state has changed to Running.

FIGURE 487 Service Status: Running Service

Service Status

REST Port:

Services Director Service Status: **running**

6. To confirm access to vTMs, click the **Services** menu, and then click **Services Director: vTM Instances**. The **vTM Instances** page appears. This page will include all of your vTMs.


Creating Services Director Reports

- [Viewing Reports and Diagnostics](#) 423
- [Viewing Logs and Generating System Dumps](#) 430
- [Working with Metering Logs](#) 432

Viewing Reports and Diagnostics

The **Home** page of the Services Director VA shows a number of summary graphs:

FIGURE 488 Home Page Summary Graphs



The **Activity** menu in the Services Director VA enables you to generate detailed reports about your current Virtual Traffic Manager (vTM) instances, bandwidth allocation, CPU utilization, and throughput. You can view how your resources are utilized so that you can adjust and reallocate resources as needed.

You can view the following reports:

- **vTM Instance Allocation** - The number of vTM instances, grouped by either instance host or feature pack, and the current status of each: *Active*, *Idle*, or *Failed*. For details, see [“The vTM Instance Allocation Report” on page 424](#).
- **Bandwidth Allocation** - The current bandwidth allocation by SKU or feature pack. For details, see [“The Bandwidth Allocation Report” on page 426](#).
- **CPU Utilization** - The current CPU utilization, grouped by either vTM instance or instance host. For details, see [“The CPU Utilization Report” on page 428](#).
- **Throughput Utilization** - The current data throughput, grouped by either vTM instance or instance host. For details, see [“The Throughput Utilization Report” on page 429](#).

Note: Historical reports are not available in this release.

The vTM Instance Allocation Report

The **vTM Instance Allocation** report summarizes the status of all instances as a series of pie charts. The main page is a two-layer pie chart. The inner layer is divided by feature pack by default, while the outer layer is divided by instance status.

The **vTM Instance Allocation** report answers these questions:

- What is the current status of my instance hosts?
- What is the current status of a particular instance host?
- What is the current status of my feature packs?
- What is the current status of a particular feature pack?

The report displays the number of instances and the status with that feature pack. You can drill down into each individual feature pack and another pie chart is presented that gives you a report on that feature pack. You also have the option to divide the inner layer of pie chart in the main page by instance host. Similarly, you can drill down into each instance host.

The **vTM Instance Allocation** report displays the current status of instances in a color coded format.

Instance Status	Color	Description
Active	Green	An instance that is currently running.
Failed	Red	An instance has failed to start.
Idle	Blue	An instance that has been deployed but is not currently running.

Pause the pointer over a specific area of the pie chart to view the feature pack or instance name (depending on the option chosen) and the number of instances.


Drill down into data by clicking an inner section of the graph.

Viewing the vTM Instance Allocation Report

To view the vTM Instance Allocation report:

1. Click **Activity > vTM Instance Allocation** to display the **vTM Instance Allocation** report page.

FIGURE 489 vTM Instance Allocation Report



2. Use the Options to change the report type:
 - **Instance host.** Then, select a specific instance host for the report, or select **All**.
 - **Feature pack.** Then, select a specific feature pack for the report, or select **All**.

When you select All, you can double-click an instance or feature pack in the pie chart to view details for the selected instance or feature pack.

3. Drill down into data by clicking one of the inner sections.

The Bandwidth Allocation Report

The **Bandwidth Allocation** report displays allocated bandwidth for your vTM instances by SKU or feature pack. When you create an instance, you must specify which feature pack you want to use; you do not specify the SKU.

The **Bandwidth Allocation** report answers these questions:

- How much bandwidth is allocated to a SKU or instance?
- How much bandwidth is unallocated for a SKU or instance?

The **Bandwidth Allocation** report is a set of pie charts. The main page is a two-layer pie chart. The inner layer is divided by licensed tied SKUs. The outer layer shows the bandwidth allocated to each of instances and total size of available bandwidth of each SKU.

Note: You cannot specify how much bandwidth you want to reserve for a given feature pack.

You can use the **Bandwidth Allocation** report to evaluate whether or not you need to reallocate bandwidth or purchase additional bandwidth licenses.


Pause the pointer over a specific area of the pie chart to view the allocated and unallocated bandwidth for a Pulse Secure Virtual Traffic Manager SKU or instance.

Drill down into data by clicking an inner section of the graph.

To view the Bandwidth Allocation report:

1. Click **Activity > Bandwidth Allocation** to display the **Bandwidth Allocation** report page.

FIGURE 490 Bandwidth Allocation Report




2. Pause over a graph section with the pointer to view a summary.

3. To drill down into a particular SKU, double-click the area you want to view. A three-layer pie chart appears:

- The inner layer displays the particular SKU.
- The middle layer is divided by feature pack created for that SKU.
- The outer layer represents the bandwidth allocated for each of instance.

FIGURE 491 Detailed Bandwidth Allocation Report

Bandwidth Allocation (ENT-ENTERPRISE 3250 Mbps/5000 Mbps Used)



The CPU Utilization Report

The **CPU Utilization** report displays real-time CPU usage by percentage over time of each instance in an *Active* state and the aggregated CPU usage of all *Active* instances on each host.

The **CPU Utilization** report answers these questions:

- How much of the CPU is being used?
- What is the average and peak percentage of the CPU being used?


The **CPU Utilization** report is a streaming line chart. The hosts and *Active* instances are listed at the bottom of line charts. You can choose which host or instance CPU usage to displayed in the chart. If you have too many *Active* instances, there is a **Filter** box from which you can filter the report by instance name. Pulse Secure recommends you use the *regular expression* name.

Viewing the CPU Utilization Report

1. Click **Activity > CPU Utilization** to display the **CPU Utilization** report page.

FIGURE 492 CPU Utilization Report

CPU Utilization



2. To view a graph of data points over time, keep the page open. Data points are graphed every ten seconds.
3. To toggle on and off the graph for an instance host, click the instance hostname at the bottom of the page.
4. To view the CPU utilization for a particular instance, enter the vTM instance name and click **Filter**.
5. To clear the data, refresh the page.

The Throughput Utilization Report

The **Throughput Utilization** report displays the real-time throughput utilization (in B/s) of each instance in an *Active* state and aggregated throughput utilization on each host.

The **Throughput Utilization** report answers these questions:

- What was the average throughput?
- What was the peak throughput?

The **Throughput Utilization** report is a streaming line chart. The real-time throughput per second and peak throughput in last hour is displayed in the chart.

The displayed throughput includes both incoming and outgoing throughput.

Review the **Throughput Utilization** report to find out which instances use the most throughput, and then compare the results to the results you expected. For example, you might expect a lot of throughput for an instance that hosts a popular site. However, if an instance is using more throughput than expected, you can try to discover why so that you can make the appropriate adjustments.


You can also use the **Throughput Utilization** report to monitor how close you are to reaching your license limitations, so that you can evaluate whether or not you should purchase additional licenses.

Pause the pointer over a specific data point to see what its value and exact time stamp were in relation to peaks.

To view the Throughput Utilization report:

1. Click **Activity > Throughput Utilization** to display the **Throughput Utilization** report page.

FIGURE 493 Throughput Utilization Report



2. To view the throughput for a particular instance, enter the vTM instance name and click **Filter**.

Viewing Logs and Generating System Dumps

You can view system logs and generate system dumps from the **Diagnose** tab.

- “Viewing System Logs” on page 430.
- “Generating System Dumps” on page 431.


Viewing System Logs

You can view current logs for the Services Director in the **System Logs** page.

1. Click **Diagnose > System Logs** to display the **System Logs** page.

FIGURE 494 System Logs Page

System Logs



2. Click << (first), < (previous), > (next) or >> (last) to navigate through the log pages.

Alternatively, type a number in the **Page** text box and click **Go** to navigate to a specific page.

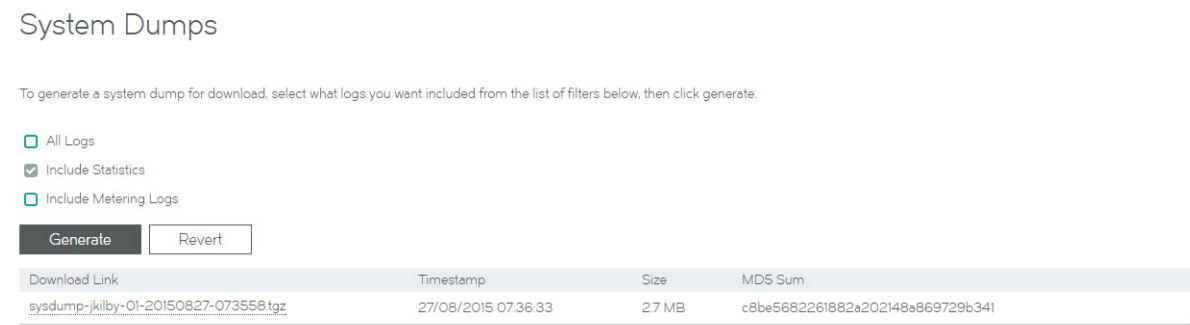
Generating System Dumps

You can generate system dumps for the Services Director from the **Diagnose** menu.

You can tailor the contents of the system dumps to include statistics if required.

1. Click **Diagnose > System Dumps** to display the **System Dumps** page.

FIGURE 495 System Dumps Page



2. Complete the configuration according to this table.

Control	Description
All Logs	Select to generate all current system logs.
Include Statistics	Select to include all statistics in system dump files.
Include Metering	Select to include all metering logs in system dump files.

3. Click **Generate** to create the system dump.

Generated logs are listed in a table of download hyperlinks.

Working with Metering Logs

The **Metering Logs** page enables you to download and manage metering log files. The files are created as .ZIP files and listed in a table. A maximum of ten metering logs can be generated by this process.

Note: Cloud Service Provider customers must ensure that SNMP is enabled on all externally-deployed vTMs to support metering.

FIGURE 496 Metering Logs Page

Metering Logs

Metering Logs Phone Home

☐ Enable Metering Logs Phone Home

Apply

Recent Metering Logs

You can download recent metering logs that have yet to be archived and phoned home to Pulse Secure. These metering logs will still be archived and phoned home at the next scheduled phone home event.

[Download](#)

Archived Metering Logs

You can download metering logs that have been archived and phoned home to Pulse Secure. If the automatic phone home process has failed, you may also retry phone home for a given metering log archive.

Download Link	Timestamp	Size	MD5 Sum	Action
services-director-metering-10_62_167_104-20180201T000010-201801.zip	01/02/2018 00:00:10	565.28 KB	7471b20de5efc46fa2677d0e285bac9a	Phone Home
services-director-metering-10_62_167_104-20180201T105810-201801.zip	01/02/2018 10:58:10	815.44 KB	deafa6db08ff22dd3c5c5eca1176db3e	Phone Home

You can monitor the capacity of the */data* partition that is used for the storage of metering logs. See **“Monitoring the Storage Capacity of Metering Logs” on page 433**.

You can also download any listed log files directly from the table.

You can also enable/disable the phone home feature from this page, see **“Configuring the Phone Home Function” on page 435**. For details of the phone home feature, refer to the *Pulse Services Director Advanced User Guide*.

Generating Metering Logs

1. Click **Diagnose > Metering Logs** to display the **Metering Logs** page.
2. Clear the **Enable Metering Logs Phone Home** check box.
3. Click **Generate** to create the metering logs.

Downloading Metering Logs


1. Click **Diagnose > Metering Logs** to display the **Metering Logs** page.
2. Click the name of the required log file (.ZIP) in the metering log table.
3. Select a save location and click **Save**.

Deleting Metering Logs

You can delete individual metering logs from the **Metering Logs** page.

1. Click **Diagnose > Metering Logs** to display the **Metering Logs** page.
2. Under **Archived Metering Logs**, locate a metering log file (.ZIP) that you want to delete.
3. To the right of the metering log entry, click the **X** control. A confirmation control appears.

FIGURE 497 The Delete Confirmation Control



4. Click **Delete**.
5. Repeat this procedure to delete additional metering logs.

Note: You can also delete all metering logs as a single action from the **Data Storage Status** page, see [“Monitoring the Storage Capacity of Metering Logs” on page 433](#).

Monitoring the Storage Capacity of Metering Logs

The Services Director VA uses its */data* partition to store both database replication logs and instance metering logs. The **Data Storage Status** page enables you to:

- View the available and used space in the */data* partition.
- Delete any archived metering logs.

Note: You cannot delete replication logs from the Services Director VA.

- Configure the number of days for which database replication logs are kept.

Note: A *Standby* Services Director can only remain offline for this many days, after which it will be unable to restore itself to the current state of the database.

To access and use the **Data Storage Status** page:

1. Click **Diagnose > Data Storage Status** to display the **Data Storage Status** page.

FIGURE 498 Data Storage Status

Data Storage Status

The Services Director VA uses its `/data` partition to store database replication logs and instance metering logs.

Use this page to:

- View the available and used space on `/data`
- Delete any archived metering logs.
- Configure the number of days for which database replication logs are kept. A standby Services Director can only remain offline for this many days, after which it will be unable to restore itself to the current state of the database.

Available space: 7.2G

Used space

Metering logs:	11M
	Delete archived logs.
Database replication logs:	8.2M

Settings

Days to keep replication logs:

2. Examine the displayed capacity information and evaluate if any action is required.
3. (Optional) Delete all metering logs. To do this:
 - a. Click **Delete archived logs** to clear replication logs from the `/data` partition. A warning dialog appears:

FIGURE 499 Deleting Replication Logs: Warning

WARNING

Only delete your archived metering logs after backing them up:
[Download metering logs](#)

Copies of monthly-archived metering logs are also downloadable from the [Metering Logs](#) page.

Please type 'delete logs' to continue:

- b. (Optional) Click **Download metering logs** to download a .TGZ file containing all metering logs.
 - c. (Optional) Click **Metering logs** to view the **Metering Logs** page. From here you can download or delete individual metering log files. See [“Downloading Metering Logs” on page 432](#) and [“Deleting Metering Logs” on page 433](#).
 - d. Type “delete logs” into the text box and click **Confirm** to delete all archived metering logs.
4. (Optional) Click the list for **Days to keep replication logs** and choose a number of days to retain replication logs, and click **Apply**.

Note: A *Standby* Services Director can only remain offline for this many days, after which it will be unable to restore itself to the current state of the database.

Configuring the Phone Home Function

You can configure the phone home function from the **Metering Logs** page.

1. Click **Diagnose > Metering Logs** to display the **Metering Logs** page.
2. Select the **Enable Metering Logs Phone Home** check box.
3. Click **Apply** to confirm your setting.

A warning e-mail will be sent every 24 hours if the phone home feature is enabled and Services Director is unable to connect to the phone home server.

Note: You can also manually phone home an individual metering log file, see [“Manually Phoning Home a Metering Log File” on page 435](#).

Manually Phoning Home a Metering Log File

You can manually phone home an individual metering log file from the **Metering Logs** page.

1. Click **Diagnose > Metering Logs** to display the **Metering Logs** page.
2. Under **Archived Metering Logs**, locate a metering log file (.ZIP) that you want to phone home.
3. To the right of the metering log entry, click the **Phone Home** action. The log file is phoned home.

Understanding Metering Logs

The Services Director automatically meters usage on a regular basis, and it optionally sends this information to Pulse Secure for billing purposes. By default, it records this information once per hour.

If a vTM instance is *Active*, the Services Director polls it to obtain total throughput and peak activity metrics. The Services Director creates a metrics log file with one line of metrics data for each vTM instance. Each line of metrics data records the name of the instance, the time elapsed since the resource was created, and the polled metrics. If an instance is not active, only the elapsed time is recorded.

If you want to generate usage or billing information, typically you process all metering log files and aggregate the results. You should use caution when aggregating data results for billing since metering records include failed deployments.

Note: Generating log files has a cumulative impact on disk space.

The Services Director collects metering data from vTM instances as follows:

- Instances that are at version 9.4 or earlier (or have no REST API enabled) have their metering collected through SNMP.
- Instances that are at version 9.5 or later with the REST API enabled have their metering collected through their REST API. If REST-based metering fails (or is not possible), the Services Director falls back to collecting using SNMP. Any metering issues will be included in the warning logs, as before.

The Services Director records the most recent metrics information for each instance in the inventory database. You can obtain this data using the REST API. The REST API does not supply bulk metrics data.

The Metering Log file is structured as follows:

- The first row contains version data for the metering log format. This first line can be ignored by customers. Ignore this line when you aggregate data for billing.
- Each subsequent row records one set of metrics for a vTM instance, in comma-separated value (CSV) format.
- The final line contains an MD5 hash of the previous lines. Ignore this line when you aggregate data for billing.

Each line of metrics contains the following fields:

Field	Description
Timestamp	The date and time, in UTC format, that the line was written.
Instance ID	The unique instance ID for the vTM instance.
Instance Tag	This information may be empty but it is included, even if empty.
Owner	(Optional) The owner of the vTM instance.
Cluster ID	The cluster for the vTM instance.
Management IP	The management IP address of the vTM instance.
Instance SKU	The SKU (or SKU combination) assigned to the vTM instance (at the time of writing to the log). The SKU might vary between readings, and variations are not recorded in the metrics log file. This property includes a hash of features applicable to the SKU. Ignore these features for billing purposes.
Feature Pack	The feature pack assigned to the vTM instance (at the time of writing to the log).
Deploy Time	The length of time (in days, hours and minutes) since the instance was deployed.

Field	Description
Throughput	<p>The number of bytes sent by the vTM instance, as recorded in the SNMP counter.</p> <p>This number is cumulative and is reset whenever the vTM instance is restarted. It is not the throughput since the latest metering action.</p> <p>To generate usage or billing information based on throughput, you should set your aggregating script to detect a drop in throughput and designate this as a restart.</p> <p>This property is applicable to active vTM instances only.</p> <ul style="list-style-type: none"> For Idle or Inactive instances, it contains a value of 0 (zero) in the log. For uncontactable instances, it contains a value of -1 in the log.
Peak Throughput	<p>The highest number of bytes sent by the vTM instance in any second of the previous hour.</p> <p>This property is applicable to active vTM instances only.</p> <ul style="list-style-type: none"> For Idle or Inactive instances, it contains a value of 0 (zero) in the log. For uncontactable instances, it contains a value of -1 in the log.
Peak Requests	<p>The highest number of requests received by the vTM instance in any second of the previous hour.</p> <p>This property is applicable to active vTM instances only.</p> <ul style="list-style-type: none"> For Idle or Inactive instances, it contains a value of 0 (zero) in the log. For uncontactable instances, it contains a value of -1 in the log.
Peak SSL Requests	<p>The highest number of Secure Socket Layer (SSL) requests received by the vTM instance in any second of the previous hour.</p> <p>This property is applicable to active vTM instances only.</p> <ul style="list-style-type: none"> For Idle or Inactive instances, it contains a value of 0 (zero) in the log. For uncontactable instances, it contains a value of -1 in the log.
Instance Bandwidth	The bandwidth (in Mbps) allocated to the vTM instance.
Record Hash	An MD5 or similar hash of the record from the Services Director license file for tamper detection. Ignore this for billing purposes.

If metrics are not collected for a period of time, peaks for the missing time are not recorded. If you reduce the metering interval, the peak values are still relative to the previous hour rather than the time since metrics were last collected.

