ExtremeManagement™

Troubleshooting Extreme Fabric Orchestrator

Release 1.2
NN48100-702
Issue 02.01
December 2017

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Extreme Networks, Inc. assumes no liability for any errors. Extreme Networks, Inc. reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

"Documentation" means information published in varying mediums which may include product information, operating instructions and performance specifications that are generally made available to users of products. Documentation does not include marketing materials. Extreme Networks shall not be responsible for any modifications, additions, or deletions to the original published version of Documentation unless such modifications, additions, or deletions were performed by or on the express behalf of Extreme Networks. End User agrees to indemnify and hold harmless Extreme Networks, Extreme Networks' agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Extreme Networks is not responsible for the contents or reliability of any linked websites referenced within this site or Documentation provided by Extreme Networks. Extreme Networks is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Extreme Networks does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Extreme Networks provides a limited warranty on Extreme Networks hardware and software. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Extreme Networks' standard warranty language, as well as information regarding support for this product while under warranty is available to Extreme Networks customers and other parties through the Extreme Networks Support website: http://www.extremenetworks.com/support under the link ""Policies" or such successor site as designated by Extreme Networks. Please note that if You acquired the product(s) from an authorized Extreme Networks Channel Partner outside of the United States and Canada, the warranty is provided to You by said Extreme Networks Channel Partner and not by Extreme Networks.

"Hosted Service" means an Extreme Networks hosted service subscription that You acquire from either Extreme Networks or an authorized Extreme Networks Channel Partner (as applicable) and which is described further in Hosted SAS or other service description documentation regarding the applicable hosted service. If You purchase a Hosted Service subscription, the foregoing limited warranty may not apply but You may be entitled to support services in connection with the Hosted Service as described further in your service description documents for the applicable Hosted Service. Contact Extreme Networks or Extreme Networks Channel Partner (as applicable) for more information.

Hosted Service

THE FOLLOWING APPLIES ONLY IF YOU PURCHASE AN EXTREME NETWORKS HOSTED SERVICE SUBSCRIPTION FROM EXTREME NETWORKS OR AN EXTREME NETWORKS CHANNEL PARTNER (AS APPLICABLE), THE TERMS OF USE FOR HOSTED SERVICES ARE AVAILABLE ON THE EXTREME NETWORKS WEBSITE, https://extremeportal.force.com OR SUCH SUCCESSOR SITE AS DESIGNATED BY EXTREME NETWORKS, AND ARE APPLICABLE TO ANYONE WHO ACCESSES OR USES THE HOSTED SERVICE. BY ACCESSING OR USING THE HOSTED SERVICE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE DOING SO (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THE TERMS OF USE. IF YOU ARE ACCEPTING THE TERMS OF USE ON BEHALF A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT THAT YOU HAVE THE AUTHORITY TO BIND SUCH ENTITY TO THESE TERMS OF USE. IF YOU DO NOT HAVE SUCH AUTHORITY, OR IF YOU DO NOT WISH TO ACCEPT THESE TERMS OF USE, YOU MUST NOT ACCESS OR USE THE HOSTED SERVICE OR AUTHORIZE ANYONE TO ACCESS OR USE THE HOSTED SERVICE.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE EXTREME NETWORKS WEBSITE, https://extremeportal.force.com OR SUCH SUCCESSOR SITE AS DESIGNATED BY EXTREME NETWORKS, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS EXTREME NETWORKS SOFTWARE, PURCHASED FROM EXTREME NETWORKS INC., ANY EXTREME NETWORKS AFFILIATE, OR AN EXTREME NETWORKS CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH EXTREME NETWORKS OR AN EXTREME NETWORKS CHANNEL PARTNER. UNLESS OTHERWISE AGREED TO BY EXTREME NETWORKS IN WRITING, EXTREME NETWORKS DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN EXTREME NETWORKS, AN EXTREME NETWORKS AFFILIATE OR AN EXTREME NETWORKS CHANNEL PARTNER; EXTREME NETWORKS RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND EXTREME NETWORKS INC. OR THE APPLICABLE EXTREME NETWORKS AFFILIATE ("EXTREME NETWORKS").

Extreme Networks grants You a license within the scope of the license types described below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License as set forth below in the Designated System(s) License (DS) section as applicable. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. "Software" means computer programs in object code, provided by Extreme Networks or an Extreme Networks Channel Partner, whether as stand-alone products, pre-installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. "Designated Processor" means a single stand-alone computing device. "Server" means a set of Designated Processors that hosts (physically or virtually) a software application to be accessed by multiple users. "Instance" means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine ("VM") or similar deployment.

License type(s)

Designated System(s) License (DS). End User may install and use each copy or an Instance of the Software only: 1) on a number of Designated Processors up to the number indicated in the order; or 2) up to the number of Instances of the Software as indicated in the order, Documentation, or as authorized by Extreme Networks in writing. Extreme Networks may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Extreme Networks through electronic means established by Extreme Networks specifically for this purpose.

Concurrent User License (CU). End User may install and use the Software on multiple Designated Processors or one or more Servers, so long as only the licensed number of Units are accessing and using the Software at any given time. A "Unit" means the unit on which Extreme Networks, at its sole discretion, bases the pricing of its licenses and can be, without limitation, an agent, port or user, an e-mail or voice mail account in the name of a person or corporate function (e.g., webmaster or helpdesk), or a directory entry in the administrative database utilized by the Software that permits one user to interface with the Software. Units may be linked to a specific, identified Server or an Instance of the Software.

CPU License (CP). End User may install and use each copy or Instance of the Software on a number of Servers up to the number indicated in the order provided that the performance capacity of the Server(s) does not exceed the performance capacity specified for the Software. End User may not re-install or operate the Software on Server(s) with a larger performance capacity without Extreme Networks' prior consent and payment of an upgrade fee.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Extreme Networks. All content on this site, the documentation, Hosted Service, and the product provided by Extreme Networks including the selection, arrangement and design of the content is owned either by Extreme Networks or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Extreme Networks. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Extreme Networks can be a criminal, as well as a civil offense under the applicable law.

Virtualization

The following applies if the product is deployed on a virtual machine. Each product has its own ordering code and license types. Note, unless otherwise stated, that each Instance of a product must be separately licensed and ordered. For example, if the end user customer or Extreme Networks Channel Partner would like to install two Instances of the same type of products, then two products of that type must be ordered.

Third Party Components

"Third Party Components" mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements ("Third Party Components"), which contain terms regarding the rights to use certain portions of the Software ("Third Party Terms"). As required, information regarding distributed Linux OS source code (for those products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the products, Documentation or on Extreme Networks' website at:http://www.extremenetworks.com/support/policies/software-licensing or such successor site as designated by Extreme Networks. The open source software license terms provided as Third Party Terms are consistent with the license rights granted in these Software License Terms, and may contain additional rights benefiting You, such as modification and distribution of the open source software. The Third Party Terms shall take precedence over these Software License Terms, solely with respect to the applicable Third Party Components to the extent that these Software License Terms impose greater restrictions on You than the applicable Third Party Terms.

The following applies only if the H.264 (AVC) codec is distributed with the product. THIS PRODUCT IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO (i) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD ("AVC VIDEO") AND/OR (ii) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE HTTP://WWW.MPEGLA.COM.

Service Provider

THE FOLLOWING APPLIES TO EXTREME NETWORKS CHANNEL PARTNER'S HOSTING OF EXTREME NETWORKS PRODUCTS OR SERVICES. THE PRODUCT OR HOSTED SERVICE MAY USE THIRD PARTY COMPONENTS SUBJECT TO THIRD PARTY TERMS AND REQUIRE A SERVICE PROVIDER TO BE INDEPENDENTLY LICENSED DIRECTLY FROM THE THIRD PARTY SUPPLIER. AN EXTREME NETWORKS CHANNEL PARTNER'S HOSTING OF EXTREME NETWORKS PRODUCTS MUST BE AUTHORIZED IN WRITING BY EXTREME NETWORKS AND IF THOSE HOSTED PRODUCTS USE OR EMBED CERTAIN THIRD PARTY SOFTWARE, INCLUDING BUT NOT LIMITED TO MICROSOFT SOFTWARE OR CODECS, THE EXTREME NETWORKS CHANNEL PARTNER IS REQUIRED TO INDEPENDENTLY OBTAIN ANY APPLICABLE LICENSE AGREEMENTS, AT THE EXTREME NETWORKS CHANNEL PARTNER'S EXPENSE, DIRECTLY FROM THE APPLICABLE THIRD PARTY SUPPLIER.

WITH RESPECT TO CODECS, IF THE EXTREME NETWORKS CHANNEL PARTNER IS HOSTING ANY PRODUCTS THAT USE OR EMBED THE G.729 CODEC, H.264 CODEC, OR H.265 CODEC, THE EXTREME NETWORKS CHANNEL PARTNER ACKNOWLEDGES AND AGREES THE EXTREME NETWORKS CHANNEL PARTNER IS RESPONSIBLE FOR ANY AND ALL RELATED FEES AND/OR ROYALTIES. THE G.729 CODEC IS LICENSED BY SIPRO LAB TELECOM INC. SEE WWW.SIPRO.COM/CONTACT.HTML. THE H.264 (AVC) CODEC IS LICENSED UNDER THE AVC PATENT PORTFOLIO LICENSE FOR THE PERSONAL USE OF A CONSUMER OR OTHER USES IN WHICH IT DOES NOT RECEIVE REMUNERATION TO: (I) ENCODE VIDEO IN COMPLIANCE WITH THE AVC STANDARD ("AVC VIDEO") AND/OR (II) DECODE AVC VIDEO THAT WAS ENCODED BY A CONSUMER ENGAGED IN A PERSONAL ACTIVITY AND/OR WAS OBTAINED FROM A VIDEO PROVIDER LICENSED TO PROVIDE AVC VIDEO. NO LICENSE IS GRANTED OR SHALL BE IMPLIED FOR ANY OTHER USE. ADDITIONAL INFORMATION FOR H.264 (AVC) AND H.265 (HEVC) CODECS MAY BE OBTAINED FROM MPEG LA, L.L.C. SEE HTTP://WWW.MPEGLA.COM.

Compliance with Laws

You acknowledge and agree that it is Your responsibility for complying with any applicable laws and regulations, including, but not limited to laws and regulations related to call recording, data privacy, intellectual property, trade secret, fraud, and music performance rights, in the country or territory where the Extreme Networks product is used.

Preventing Toll Fraud

"Toll Fraud" is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Security Vulnerabilities

Information about Extreme Networks' security support policies can be found in the Global Technical Assistance Center Knowledgebase at https://gtacknowledge.extremenetworks.com/.

Downloading Documentation

For the most current versions of Documentation, see the Extreme Networks Support website: http://documentation.extremenetworks.com, or such successor site as designated by Extreme Networks.

Contact Extreme Networks Support

See the Extreme Networks Support website:http://www.extremenetworks.com/support for product or Hosted Service notices and articles, or to report a problem with your Extreme Networks product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Extreme Networks Support website:http://www.extremenetworks.com/support/contact/ (or such successor site as designated by Extreme Networks), scroll to the bottom of the page, and select Contact Extreme Networks Support.

Contact Avaya Support

See the Avaya Support website: https://support.avaya.com for product or Hosted Service notices and articles, or to report a problem with your Avaya product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: https://support.avaya.com (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Trademarks

The trademarks, logos and service marks (“Marks”) displayed in this site, the Documentation, Hosted Service(s), and product(s) provided by Extreme Networks are the registered or unregistered Marks of Extreme Networks, Inc., its affiliates, its licensors, its suppliers, or other third parties. Users are not permitted to use such Marks without prior written consent from Extreme Networks or such third party which may own the Mark. Nothing contained in this site, the Documentation, Hosted Service(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Extreme Networks or the applicable third party.

Extreme Networks is a registered trademark of Extreme Networks, Inc.

All non-Extreme Networks trademarks are the property of their respective owners. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

For additional information on Extreme Networks trademarks, please see: http://www.extremenetworks.com/company/legal/

Contents

Chapter 1: Preface

Purpose

The current document describes how to use troubleshooting tools and utilities for Extreme Fabric Orchestrator (EFO).

Training

Ongoing product training is available. For more information or to register, you can access the Web site at www.extremenetworks.com/education/.

Providing Feedback to Us

We are always striving to improve our documentation and help you work better, so we want to hear from you! We welcome all feedback but especially want to know about:

- Content errors or confusing or conflicting information.
- Ideas for improvements to our documentation so you can find the information you need faster.
- Broken links or usability issues.

If you would like to provide feedback to the Extreme Networks Information Development team about this document, please contact us using our short online feedback form. You can also email us directly at internalinfodev@extremenetworks.com

Getting Help

Product purchased from Extreme Networks

If you purchased your product from Extreme Networks, use the following support contact information to get help.

If you require assistance, contact Extreme Networks using one of the following methods:

- GTAC (Global Technical Assistance Center) for Immediate Support
  - Phone: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact
  - Email: support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- GTAC Knowledge – Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- The Hub – A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- Support Portal – Manage cases, downloads, service contracts, product licensing, and training and certifications.

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Product purchased from Avaya

If you purchased your product from Avaya, use the following support contact information to get help.

Go to the Avaya Support website at http://support.avaya.com for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes, downloads, and resolutions to issues. Use the online service request system to create a service request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

Extreme Networks Documentation

To find Extreme Networks product guides, visit our documentation pages at:

Current Product Documentation	www.extremenetworks.com/documentation/
Archived Documentation (for previous versions and legacy products)	www.extremenetworks.com/support/documentation-archives/
Release Notes	www.extremenetworks.com/support/release-notes

Open Source Declarations

Some software files have been licensed under certain open source licenses. More information is available at: www.extremenetworks.com/support/policies/software-licensing.

Subscribing to service notifications

Subscribe to receive an email notification for product and software release announcements, Vulnerability Notices, and Service Notifications.

About this task

You can modify your product selections at any time.

Procedure

1. In an Internet browser, go to http://www.extremenetworks.com/support/service-notification-form/ .
2. Type your first and last name.
3. Type the name of your company.
4. Type your email address.
5. Type your job title.
6. Select the industry in which your company operates.
7. Confirm your geographic information is correct.
8. Select the products for which you would like to receive notifications.
9. Click **Submit**.

Chapter 2: New in this document

The following sections detail what’s new in *Troubleshooting Extreme Fabric Orchestrator*, NN48100–702. See *Extreme Fabric Orchestrator Release Notes* for a full list of features.

There are no feature changes.

Chapter 3: Troubleshooting deploying

Use the following information to troubleshoot deploying.

Enabling or disabling the Out-of-band network

Use the following information to enable or disable an out-of-band network. An out-of-band (OOB) network is a separate network you can enable to manage the device network, rather than using the appliance management network.

Procedure

1. Login into the KVM hypervisor using SSH as the root user.
2. Run the following command:

   ```
   #/bin/bash//opt/avaya/afo/infra/scripts/configOOB.sh
   ```

3. If an OOB network exists the system asks if you want to disable the OOB network. Select y for Yes.

Example

Note:

If an OOB network does not exist, the system asks if you want to enable an OOB network. The following example is for a disabling scenario.

```
[root@afo-server1-kvm ~]# bash /opt/avaya/afo/infra/scripts/configOOB.sh
Already OOB network is configured, do you want to disable OOB network [y/n] or [Y/N]? y
```

Chapter 4: Network Discovery and Monitoring troubleshooting

Use the following information to troubleshoot Network Discovery and Monitoring.

Default domain

The system has a Default domain that cannot be removed. You can remove the content under the Default domain.

All configuration tasks are performed only on the elements in the Default domain.

The default setting for the Default domain is Avaya-only discovery.

Avaya-only discovery means that the system only discovers Avaya devices and Avaya partner devices. The default domain discovery configuration is read-only and cannot be changed.

Ensuring proper device credentials

Condition

Ensure device and server credentials are entered before you start Network Discovery.

Solution

Select **Administration** > **Credentials** and configure the proper device credential settings.

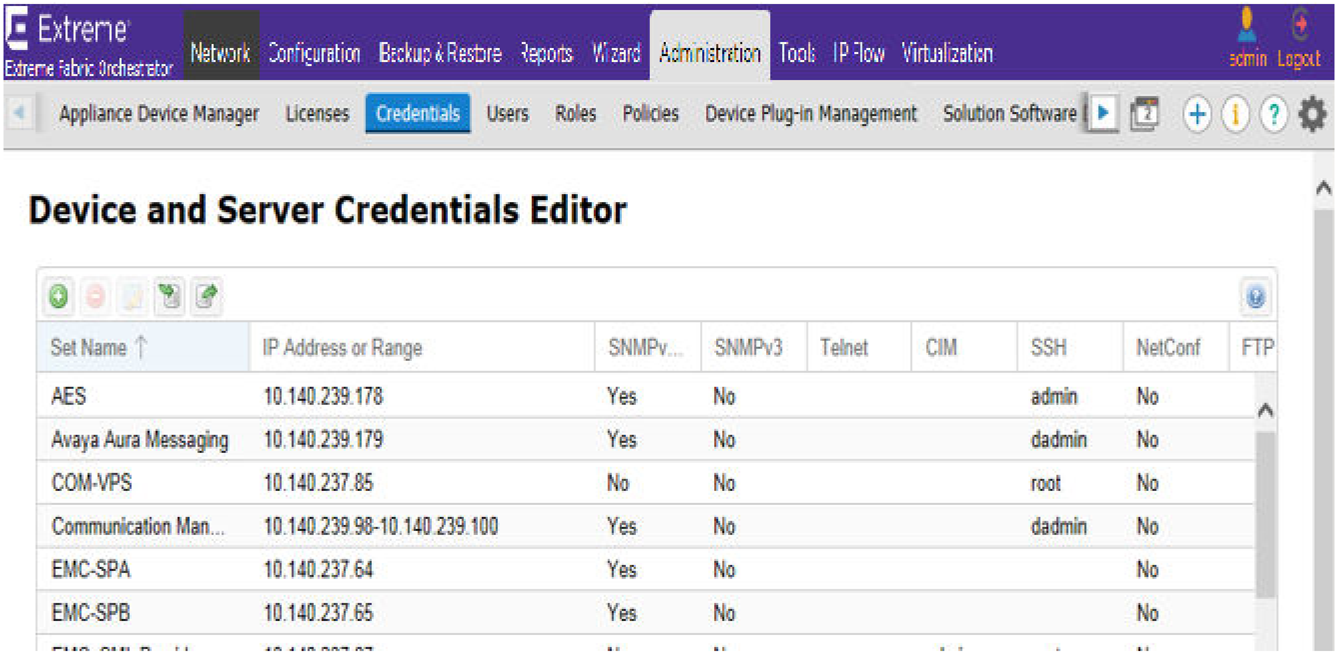

Device and Server Credentials Editor

Set Name	IP Address or Range	SNMPv...	SNMPv3	Telnet	CIM	SSH	NetConf	FTP
AES	10.140.239.178	Yes	No			admin	No	
Avaya Aura Messaging	10.140.239.179	Yes	No			dadmin	No	
COM-VPS	10.140.237.85	No	No			root	No	
Communication Man...	10.140.239.98-10.140.239.100	Yes	No			dadmin	No	
EMC-SPA	10.140.237.64	Yes	No				No	
EMC-SPB	10.140.237.65	Yes	No				No	

Rediscovery policies

As part of a rediscovery under **Network** > **Discovery**, you can apply a variety of different policies, which include:

- Retain missing equipment if possible
- Rediscover from scratch retaining states
- Rediscover from scratch
- Retain equipment unless marked to remove

By default a rediscovery adds equipment back to the domain if the system finds equipment in the network, even if the equipment is marked for removal. Use this policy to remove equipment marked for removal, even if the equipment exists in the network during the time of a rediscovery.

Discovery status indicates partial discovery

Condition

After you do a Network Discovery in **Network** > **Discovery**, the system can indicate only a partial discovery.

You can receive a message: `Discovery was partially successful. 1 seed of 3 are not discovered`.

Cause

- The system has not reached all seeds for discovery because one of the seed routers was not up, or the credentials were not provided.
- The discovery has run out of licenses.

Solution

1. To debug, go to **Reports** > **Discovery Reports** icon, and look under the category SEED.
2. Ensure that all of the see routers are up.
3. Ensure that the proper credentials are provided in **Administration** > **Credentials**.
4. Ensure enough licenses exist.
5. Try to do a new discovery.

   After a discovery is finished, under the Discovery Status Summary:

   - The Discovery State displays as Completed.
   - The Discovery Type displays as Full discovery.
   - The system displays as Discovery successful.

Discovery log message displays as excluded by discovery filter

Condition

You receive a log message that displays as: `Potential managed device <IP: 198.51.100.0> was not discovered. It’s excluded by discovery filter. Referenced by....`

Cause

This is a general warning in the log that means an IP address was found in the SONMP, LLDP, or CP table of a discovered device, but the IP object was excluded to do a limit on the discovery range.

Solution

1. Check the limits you have configured for the discovery. By default, Avaya-only devices are discovered.
2. Check the Discovery problem report. Select **Network** > **Discovery**, and then click the **Discovery Problem Report** icon.

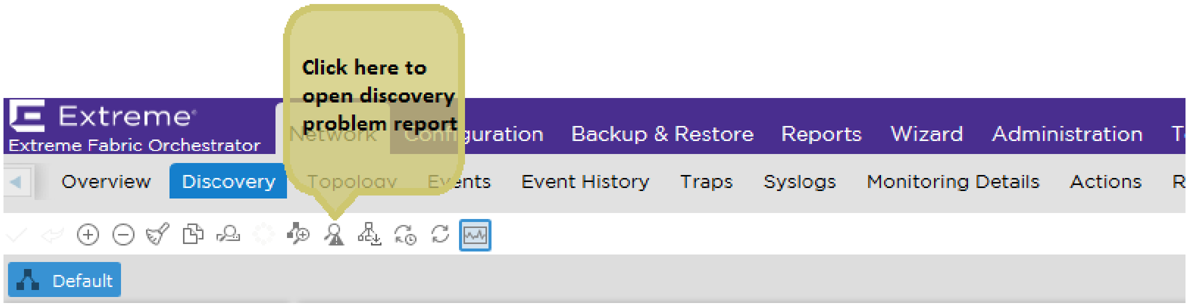

Discovery is missing devices because of no SNMP response

Condition

The system does not discover devices on the network, and the log says: `No SNMP response from suspected SNMP device — check credentials.`

Cause

This is a general warning in the log that means that an IP address was found in the forwarding table of a discovered device, but the IP object did not respond to SNMP.

Solution

1. Check the credentials for the device under **Administration** > **Credentials**.
2. Adjust the SNMP timeout in **Global Preferences**. Click on the **Preferences** icon on the top right, and then select the **Global** tab.
3. Check the Discovery problem report. Select **Network** > **Discovery**, then click the **Discovery Problem Report** icon.

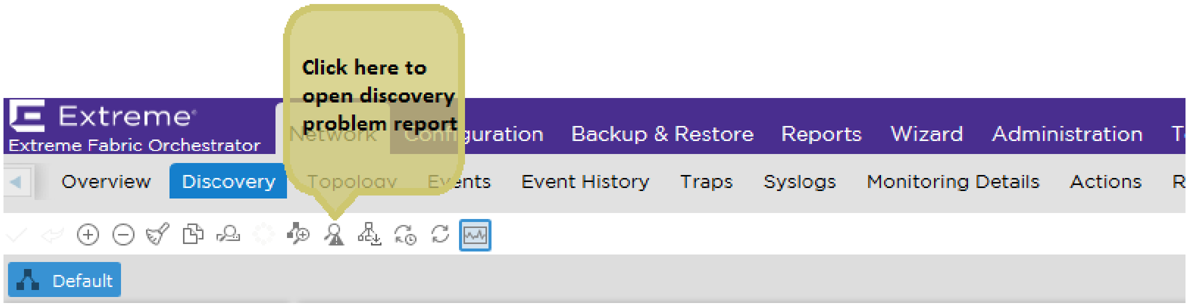

Discovery log says no response from ICMP ping and SNMP request

Condition

The system did not discover devices on the network, and the log says: `No response from ICMP ping and SNMP request...`

Cause

This is a general message in the log that an IP address was found in some other device table but the IP address is not responding to ping or SNMP, which suggests the device is now physically disconnected or no route exists to the device.

Solution

1. Use ping to query the device.
2. Go to **Tools** > **MIB Browser** and use MIB Browser to query the device.

Topology is not displayed in the Network Topology

Condition

The discovery has completed, and the topology is not displayed in the network topology.

Cause

Invalid IP address or device credentials.

Solution

1. Check if you have provided a valid device seed IP address by going to **Network** > **Discovery**.
2. Correct the device seed IP address if necessary.
3. Check if the device credentials are valid by going to **Administration** > **Credentials**.
4. Correct the device credentials if necessary.
5. Restart the discovery.

Determine when to run the network discovery

Condition

Determine when to start the network discovery.

Solution

1. You can start the network discovery after the system is up, and you have attached the devices to the Fabric through an edge switch, or a proxy.
2. Go to **Network** > **Discovery**.
3. Select the domain you want to discover.
4. Click **Discover selected domain** from the top menu bar.
5. Click **OK** to start the discovery.
6. To update the latest topology, click the **Refresh** button.
7. If you add new r devices to the Fabric, refresh the topology view within a period of five minutes to see the updated topology.

Unable to launch any Discovery or Monitoring page

Condition

Unable to launch any Discovery or Monitoring page.

Discovery and **Monitoring** pages are under the **Network** top level menu.

MIB Browser and **MIB Query** pages are under the **Tools** top level menu.

Solution

1. If you have an issue with the **Discovery**, **Monitoring**, **MIB Browser**, or **MIB Query** pages not displaying, check the status of kbmd service in the Monitoring virtual machine (VM). To check for a problem, select **Administration** > **Appliance Device Manager**. Click on the Monitoring module in the Services portlet.

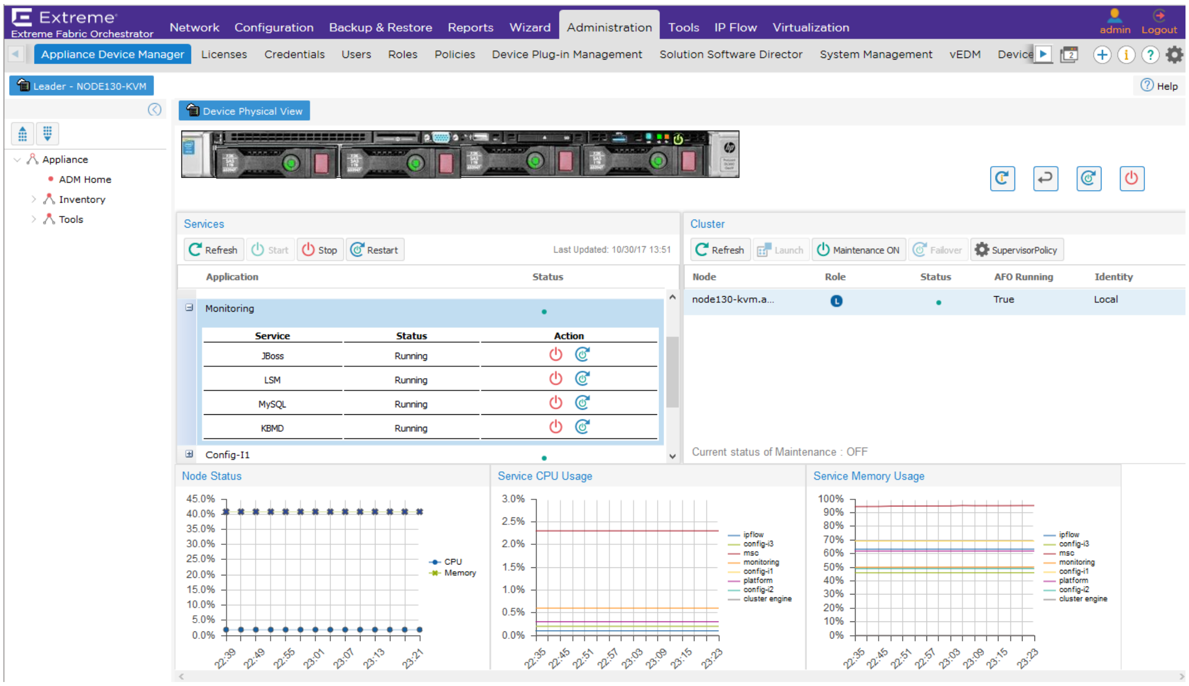
Extreme
Extreme Fabric Orchestrator
Network Configuration Backup & Restore Reports Wizard Administration Tools IP Flow Virtualization
admin Logout
Appliance Device Manager Licenses Credentials Users Roles Policies Device Plug-in Management Solution Software Director System Management vEDM Device
Leader - NODE130-KVM
Help
Appliance
ADM Home
Inventory
Tools
Device Physical View
Services
Refresh Start Stop Restart
Cluster
Refresh Launch Maintenance ON Failover SupervisorPolicy
Node Status
CPU
Memory
Service CPU Usage
Service Memory Usage
ipflow
config-i3
msc
monitoring
config-i1
platform
config-i2
cluster engine

Last Updated: 10/30/17 13:51

Application	Status
Monitoring	●
Config-I1	●

Service	Status	Action
JBoss	Running	
LSM	Running	
MySQL	Running	
KBMD	Running	

Node	Role	Status	AFO Running	Identity
node130-kvm.a...	L	●	True	Local

Current status of Maintenance : OFF

2. If you want to restart the services you can do so under **Administration** > **Appliance Device Manager** page.

3. Under Appliance Device Manager, select **Restart Service**.

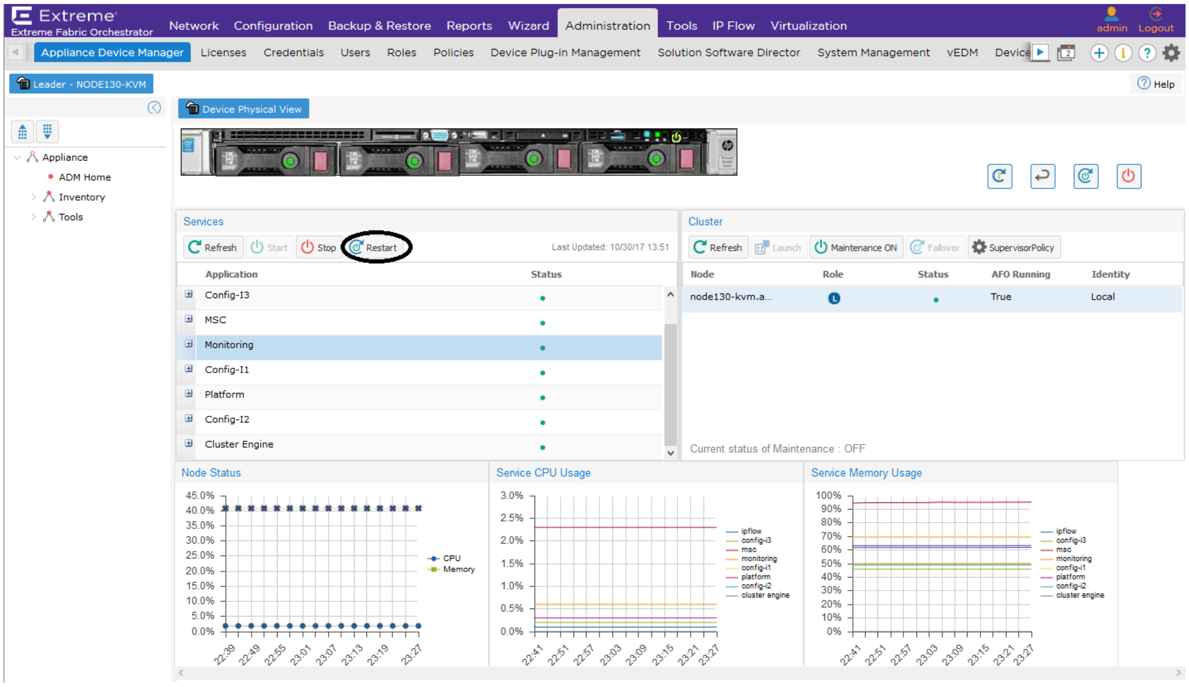
Extreme
Extreme Fabric Orchestrator
Network Configuration Backup & Restore Reports Wizard Administration Tools IP Flow Virtualization
admin Logout
Appliance Device Manager Licenses Credentials Users Roles Policies Device Plug-in Management Solution Software Director System Management vEDM Device
Leader - NODE130-KVM
Help
Appliance
ADM Home
Inventory
Tools
Device Physical View
Services
Refresh Start Stop Restart
Cluster
Refresh Launch Maintenance ON Failover SupervisorPolicy
Node Status
CPU
Memory
Service CPU Usage
Service Memory Usage
ipflow
config-i3
msc
monitoring
config-i1
platform
config-i2
cluster engine

Last Updated: 10/30/17 13:51

Application	Status
Config-I3	●
MSC	●
Monitoring	●
Config-I1	●
Platform	●
Config-I2	●
Cluster Engine	●

Node	Role	Status	AFO Running	Identity
node130-kvm.a...	L	●	True	Local

Current status of Maintenance : OFF

Discovery and Monitoring pages do not display after installing a new license

Condition

Discovery and Monitoring pages do not display after installing a new license.

Solution

After installing a new license after your trial period expires, or if you upgrade to a higher license tier, you must log in again for the change to licensing to be effective.

Stopped one agent in Monitoring Details and cannot start again

Condition

A user stopped one agent in Monitoring Details, under **Network** > **Monitoring Details**, and cannot start the agent again.

⚠ **Warning:**

- If you stop an agent, you must stop and restart monitoring for the whole domain to restart the agent.
- Do not start and stop the individual agents.

The run state of an agent indicates if the agent is running, not if the system is monitoring the agent. In order for monitoring to occur within an agent both the domain monitoring state and the agent run state must be green.

When the domain state is set to monitoring, the agents are created and started so they enter the running state and the system sends monitoring requests to the agents.

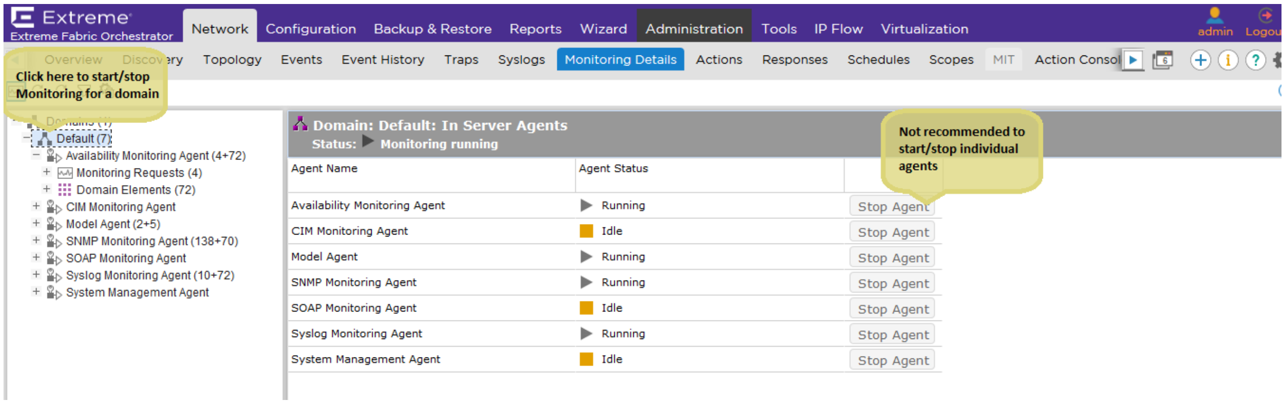

Click here to start/stop Monitoring for a domain

Domain: Default: In Server Agents
Status: ▶ Monitoring running

Not recommended to start/stop individual agents

Agent Name	Agent Status	
Availability Monitoring Agent	▶ Running	Stop Agent
CIM Monitoring Agent	■ Idle	Stop Agent
Model Agent	▶ Running	Stop Agent
SNMP Monitoring Agent	▶ Running	Stop Agent
SOAP Monitoring Agent	■ Idle	Stop Agent
Syslog Monitoring Agent	▶ Running	Stop Agent
System Management Agent	■ Idle	Stop Agent

Unable to launch SPBM L2 Diagnose Tools

Condition

SPBM Diagnose Tools fail to launch with Warning: com.rocketsoft.snmp.QueryTimoutException.

Network > **Topology** > **SPBM view**. Shift click on two devices, then right click and select **SPBM Diagnose Tools**.

Solution

You must provide the SNMP write community for all SPBM enabled devices for the SPBM L2 Diagnose test to run successfully.

Chapter 5: Troubleshooting licensing

Use the following information to troubleshoot licensing.

Inventory count in dashboard

License tiers are 250-nodes, 1500-nodes, and 5000-nodes. These counts represent the number of approved vendor switches licensed for the configuration.

Monitoring nodes include a further three times of this count, and include servers, IP phones, and other multi-vendor managed nodes. For example, a 1500-node license includes 1500 approved vendor nodes for configuration, and a total of 6000 nodes for monitoring.

The managed node count in the License Summary dashboard widget includes all managed nodes (switches, servers, and phones).

To ensure full discovery, a buffer of an extra 25 nodes is provided with all licensing tiers.

Chapter 6: CLI troubleshooting

Use the following information to troubleshoot the system using command line interface (CLI).

Command Line Interface

You can use the command line interface (CLI) to administer and use some of the key features. The CLI provides a number of commands to perform the administrative and troubleshooting tasks.

Through the CLI, you can:

- Perform a factory reset.
- View the hardware resource usage.
- Perform a health check.
- Start, stop, or restart the application service.
- Update and edit the network configuration.
- Configure the NTP.
- Update the iLO settings.
- View the HA health status.
- View the Host ID

Note:

You must use the CLI to perform all the troubleshooting related tasks for the system.

CLI commands

The following table lists the CLI commands available through CLI:

You can run the following commands with root user credentials from the KVM hypervisor.

Command	Description	Syntax
Factory reset	Allows you to re-deploy services on the server.	`cluster-factory-reset`
Resource usage	Displays the current CPU and memory usage of each virtual machine.	`cluster-resource-usage`
Health check	Allows you to check the status of the applications running on each virtual machine.	`cluster-health-check`

Service	Allows you to easily stop, start, or restart the application service.	• To view the help menu
		`cluster-service -help`
		• To stop the services in all VMs:
		`cluster-service -action stop -serviceid all`
		• To start services in all VMs:
		`cluster-service -action start -serviceid all`
		• To restart services in all VMs:
		`cluster-service -action restart -serviceid all`
		• To check the status of all VMs:
		`cluster-service -action status -serviceid all`
		• To stop applications in a particular VM:
		`cluster-service -action stop -serviceid <platform|fault|flow|config|config1|config2|msc>`
		• To start applications in a particular VM:
		`cluster-service -action start -serviceid <platform|fault|flow|config|config1|config2|msc>`
		• To restart applications in a particular VM:
		`cluster-service -action restart -serviceid <platform|fault|flow|`

Table continues…

Command	Description	Syntax
		config|config1|config2|msc> • To check the status of applications in a particular VM: cluster-service -action status -serviceid <platform|fault|flow|config|config1|config2|msc> • To configure the NTP cluster-configure-ntp • To update the iLO settings cluster-update-iloinfo
Network change	Allows you to update and change the network details post deployment. **Note:** The process takes approximately 45 minutes to complete. **Important:** If you change the IP address or FQDN, during the next login you are prompted to change the password.	cluster-network-config
Network information	Displays the hostname and IP address of all the virtual machines that are configured.	cluster-info
HostID	Provides the HostID for generating a license.	cluster-hostid
HA health status	Provides the HA health status at any given point of time.	cluster-ha-status

Factory reset

The factory reset utility allows you to re-deploy the appliance on the server. Before you perform a factory reset, ensure that you have taken a backup of the existing data.

Warning:

This command wipes all settings, configurations, and information from the system, and takes your system back to the factory settings. Make sure that you take a backup and store the backup outside of the appliance before you perform a factory reset. The factory reset utility brings the server back into the same state of the factory build, and all the existing data will be erased.

Note:

Backup the data before performing a factory reset.

Condition

If the device is in an unusable state, and you are not able to troubleshoot.

Caution:

Contact the customer support center before performing factory reset.

Solution

1. Login to Appliance base Platform (Hypervisor) as a root user.
2. Run the **`cluster-factory-reset`** command on the Command Line Interface.

   The system prompts you to confirm before the system starts the factory reset procedure.

Choice Option	Choice Description
y	Enter `y` to continue the reset
n	Enter `n` to cancel the reset

3. The system prompts you to either restart the server to begin the configuration or shutdown the server to configure later.

Choice Option	Choice Description
r	Enter `r` to restart the server to begin the configuration
s	Enter `s` to shutdown the server to configure later

```
[root@perf-cluster1-kvm ~]# cluster-factory-reset

Factory Reset will re-deploy Fabric Connect Services on this server.
Existing data will not be retained. Hence it is strongly recommended to backup
the data before proceeding further.

Do you want to continue? [y/n]:y

Are you sure you want to continue Factory Reset?[y/n]:y
>> 'Platform(Common Service)' service re-deployed successfully.
>> 'IP Flow' service re-deployed successfully.
>> 'Monitoring' service re-deployed successfully.
>> 'Configuration 1' service re-deployed successfully.
>> 'Configuration 2' service re-deployed successfully.
>> 'Configuration 3' service re-deployed successfully.
>> 'Avaya Diagnostic Server' service re-deployed successfully.
>> 'Management Server Console' service re-deployed successfully.
```

```
Factory Reset completed successfully.
r - Restart the Server
s - Shutdown the server
Choose option [r/s]:
```

Displaying the current CPU memory usage on VMs

You can use the Hardware Resource Usage utility to provide the current CPU and memory usage of each virtual machine on the system. This utility is present on the Management Server Console (MSC).

Procedure

1. Use SSH to log in to the KVM hypervisor as the root user.
2. Enter the following command to provide the current CPU and memory usage of each virtual machine on the system:

   ```
   cluster-resource-usage
   ```

Example

Use SSH to login to the KVM hypervisor as the root user and run the command to view the current CPU and memory usage:

```
login as: root
Access denied
admin@192.0.2.1's password:
Last login: Sun May 22 07:57:34 2017 from 198.51.100.0
Last login: Sun May 22 Sun May 22 07:57:34 2017 from 192.0.2.1
[root@perf-cluster1-kvm ~]# cluster-resource-usage
.--------------------------------------------------------.
|     Avaya Fabric Orchestrator Resource Usage Status    |
+----------------------------+------------+--------------+
|           Service          |   CPU (%)  |  Memory (%)  |
+----------------------------+------------+--------------+
| ADS                        | 0.1        | 33.84        |
| Config1                    | 0.1        | 49.42        |
| Config2                    | 0.1        | 48.98        |
| Config3                    | 0.1        | 48.97        |
| IPFLOW                     | 0.1        | 31.76        |
| Monitoring                 | 0.3        | 22.49        |
| MSC                        | 0.1        | 55.69        |
| Platform                   | 0.2        | 72.38        |
'----------------------------+------------+--------------'
```

Service utility

The Service utility allows you to easily stop, start, or restart the application service for any particular or multiple virtual machines without hindering your current application service.

The options available for the Service utility are:

- Help command
- Command to start, stop, or restart the application service on a particular virtual machine (VM), multiple VMs, or all VMs.
- Command to check the status of a particular VM, multiple VMs, or all VMs.

Multiple services run on different virtual machines. There may be situations where you want to start, stop, or restart the application service for any particular or multiple virtual machines.

Displaying the service ID information

The help menu displays the list of service IDs along with their descriptions.

Procedure

1. Use SSH to log into the KVM hypervisor as the root user.
2. Enter the following command to display a list of service IDs:

   `cluster-service —help`

Example

Use SSH to log into the KVM hypervisor as the root user, and enter the **cluster-service -help** to display the service ID descriptions.

```
login as: root
Access denied
Using keyboard-interactive authentication.
Password:
Last login: Sun May 22 Sun May 22 07:57:34 2017 from 198.51.100.0
[root@perf-cluster1-kvm ~]# cluster-service -help
.----------------------------------------.
|         Service ID Description         |
+------------+---------------------------+
| Service ID |        Description        |
+------------+---------------------------+
| platform   | Platform (Common Service) |
| fault      | Monitoring                |
| flow       | IP Flow                   |
| config1    | Configuration 1           |
| config2    | Configuration 2           |
| config3    | Configuration 3           |
| msc        | Management Server Console |
'------------+---------------------------'
[Example:
cluster-service -action stop -serviceid all
cluster-service -action stop -serviceid platform,config1,msc
cluster-service -action start -serviceid all
cluster-service -action start -serviceid config1,msc
cluster-service -action restart -serviceid all
cluster-service -action restart -serviceid config1,msc
cluster-service -action status -serviceid all
cluster-service -action status -serviceid msc
```

Stopping an application service

Use the following procedure to stop an application service or to stop all application services.

Procedure

1. Use SSH to log into the KVM hypervisor as the root user.
2. Enter the following command to stop an application service:

   ```
   cluster-service -action stop -serviceid <platform|fault|flow|
   config1|config2|config3|msc>
   ```

3. Enter the following command to stop all application services:

   ```
   cluster-service -action stop -serviceid all
   ```

Example

Use SSH to log into the KVM hypervisor as the root user, and enter the **`cluster-service -action stop -serviceid <platform|fault|flow|config1|config2|config3|msc>`** to stop a particular service.

```
login as: root
Access denied
admin@192.0.2.1's password:
Last login: Sun May 22 07:57:34 2017 from 198.51.100.0
[root@perf-cluster1-kvm ~]# cluster-service -action stop -serviceid flow
```

Starting an application service

Use the following procedure to start an application service on a particular VM.

About this task

After you start or restart services, the system can take five to 10 minutes before these services are completely operational.

Procedure

1. Use SSH to log into the KVM hypervisor as the root user.
2. Enter the following command to start an application service:

   ```
   cluster-service -action start -serviceid <platform|fault|flow|
   config1|config2|config3|msc>
   ```

3. Enter the following command to start all application services:

   cluster-service -action start -serviceid all

4. Enter the following command to restart an application service:

   ```
   cluster—service action restart -serviceid <platform|fault|flow|
   config1|config2|config3|msc>
   ```

5. Enter the following command to restart all application services:

   cluster-service -action restart -serviceid all

Example

Use SSH to log into the KVM hypervisor as the root user, and enter the **`cluster-service -action start -serviceid <platform|fault|flow|config1|config2|config3|msc>`** to start a particular service.

```
login as: root
Access denied
admin@192.0.2.1's password:
Last login: Sun May 22 07:57:34 2017 from 198.51.100.0
[root@perf-cluster1-kvm ~]# cluster-service -action start -serviceid flow

.--------------------------------------------------.
|     Avaya Fabric Orchestrator Service Status     |
+--------------+---------------------+-------------+
|   Service    |     Application     |   Status    |
+--------------+---------------------+-------------+
| IPFLOW       | JBoss               | Up          |
| IPFLOW       | MySQL               | Up          |
| IPFLOW       | IPFix Collector     | Up          |
'--------------+---------------------+-------------'

Note: Startup of services initiated. It can take 5-10 minutes for these services to be
completely operational.
```

Checking the status of applications running on each VM

Use the following procedure to check the status of applications running on each virtual machine (VM).

About this task

A status of up means the application is running. A status of down means the application is not running.

Procedure

1. Use SSH to log into the KVM hypervisor as the root user.
2. Use the following command to check the status of all of the applications.

   ```
   cluster-service -action status -serviceid all
   ```

3. Use the following command to check the status of a particular application.

   ```
   cluster-service -action status -serviceid <platform|fault|flow|
   config1|config2|config3|msc>
   ```

Example

Use SSH to log into the KVM hypervisor as the root user.

```
login as: root
Access denied
admin@192.0.2.1's password:
Last login: Sun May 22 07:57:34 2017 from 198.51.100.0
```

```
[root@perf-cluster1-kvm ~]# cluster-service -action status -serviceid all

.-------------------------------------------------------.
|        Avaya Fabric Orchestrator Service Status       |
+---------------------------+-----------------+---------+
|          Service          |   Application   | Status  |
+---------------------------+-----------------+---------+
| Platform                  | JBoss           | Up      |
| Platform                  | PostgreSQL      | Up      |
| Platform                  | CND             | Up      |
| Monitoring                | JBoss           | Up      |
| Monitoring                | MySQL           | Up      |
| Monitoring                | LSM             | Up      |
| Monitoring                | KBMD            | Up      |
| IPFLOW                    | JBoss           | Up      |
| IPFLOW                    | MySQL           | Up      |
| IPFLOW                    | IPFix Collector | Up      |
| Configuration 1           | JBoss           | Up      |
| Configuration 2           | JBoss           | Up      |
| Configuration 3           | JBoss           | Up      |
| Management Server Console | JBoss           | Up      |
'---------------------------+-----------------+---------'
[root@perf-cluster1-kvm ~]# cluster-service -action status -serviceid config2

.-------------------------------------------------.
|     Avaya Fabric Orchestrator Service Status    |
+--------------------+----------------+-----------+
|      Service       |  Application   |  Status   |
+--------------------+----------------+-----------+
| Configuration 2    | JBoss          | Up        |
'--------------------+----------------+-----------'
```

Performing a health check on VMs

The health check utility allows you to check the status of the applications running on each virtual machine. With the health check utility, you can discover if an application has an issue to prevent jobs from being scheduled or run on that particular application. Using health check increases the reliability and throughput of the cluster and reduces preventable job failures.

The system displays the health status of all of the applications. A status of UP means the application is running properly. A status of DOWN means the application is not running properly.

Procedure

1. Use SSH to log into the KVM hypervisor as the root user.
2. Enter the following command to check the health status of the applications running on each VM:

   ```
   cluster-health-check
   ```

Example

Use SSH to log into the KVM hypervisor as the root user, and enter the following command to check the health status of the applications running on each VM:

```
login as: root
Access denied
admin@192.0.2.1's password:
Last login: Sun May 22 07:57:34 2017 from 198.51.100.0
[root@perf-cluster1-kvm ~]# cluster-health-check

.-------------------------------------------------------.
|        Avaya Fabric Orchestrator Service Status       |
+---------------------------+-----------------+---------+
|          Service          |   Application   | Status  |
+---------------------------+-----------------+---------+
| Platform                  | JBoss           | Up      |
| Platform                  | PostgreSQL      | Up      |
| Platform                  | CND             | Up      |
| Monitoring                | JBoss           | Up      |
| Monitoring                | MySQL           | Up      |
| Monitoring                | LSM             | Up      |
| Monitoring                | KBMD            | Up      |
| IPFLOW                    | JBoss           | Up      |
| IPFLOW                    | MySQL           | Up      |
| IPFLOW                    | IPFix Collector | Up      |
| Configuration 1           | JBoss           | Up      |
| Configuration 2           | JBoss           | Up      |
| Configuration 3           | JBoss           | Up      |
| Management Server Console | JBoss           | Up      |
'---------------------------+-----------------+---------'
```

Updating network details

Use the following procedure to update network details post deployment.

Note:

The system can take approximately 45 minutes to complete the configuration.

Important:

If you change the IP address or FQDN, during the next logon you are prompted to change the password.

Procedure

1. Use SSH to log into the KVM hypervisor as the root user.
2. Use the following command to update network details post deployment.

   `cluster-network-config`

3. Enter the prefix and the domain name for the appliance for auto generating the hostname, and then enter the IP address range for configuring the applications.

Example

Use SSH to log into the KVM hypervisor as the root user.

```
login as: root
Access denied
admin@192.0.2.1's password:
Last login: Sun May 22 07:57:34 2017 from 198.51.100.0
[root@perf-cluster1-kvm ~]# cluster-network-config
Enter Prefix name for the appliance for auto generating the Hostname [e.g., cluster1]:
cluster1
Enter Domain name for the appliance for auto generating the Hostname [e.g.,
domain.com]: avaya.com
Application Network Configuration Details:
   1. Appliance base Platform (Hypervisor)
   2. Management Server Console (MSC)
   3. Platform
   4. Monitoring
   5. IPFLOW
   6. Three instance of Configuration
   7. Avaya Diagnostic Server
Enter IP Address range for configuring the above applications [Multiple IP Addresses
separated by comma]
 (e.g., [192.0.2.0-192.0.2.24] or [192.0.2.0-192.0.2.20,192.0.2.21]):
198.51.100.0-203.0.113.20
```

Displaying the hostname and IP address of VMs

Use the following procedure to display the hostname and IP address of all the virtual machines (VMs) configured.

Procedure

1. Use SSH to log into the KVM hypervisor as the root user.
2. Enter the following command to display the hostname and IP address of the VMs:

   ```
   cluster-info
   ```

Viewing the host ID for generating a license

Use this procedure to view the host ID for generating a license. You require the host ID of the WebLM server to obtain the license from PLDS.

Procedure

1. Use SSH to log in to the KVM hypervisor as the root user.
2. Use the following command to view the host ID:

   ```
   cluster-hostid
   ```

Example

Use SSH to log in to the KVM hypervisor as the root user, and use the **`cluster-hostid`** command to view the host ID.

```
login as: root
Access denied
admin@192.0.2.1's password:
Last login: Sun May 22 07:57:34 2017 from 198.51.100.0
Last login: Sun May 22 Sun May 22 07:57:34 2017 from 192.0.2.1
[root@cluster1-kvm ~]# cluster-hostid
VA872A8286F7-VC3839159F29
```

KVM hypervisor tools

Perform operations on services from the hypervisor of the system.

Shutting down the virtual machine

Use the following procedure to shutdown the virtual machine.

Procedure

1. Use SSH to log into the KVM hypervisor as the root user.
2. Use the following command to shutdown the virtual machine:

   ```
   /bin/systemctlstop afo-service.service
   ```

Starting the virtual machine

Use the following procedure to start the virtual machine.

Procedure

1. Use SSH to log into the KVM hypervisor as the root user.
2. Use the following command to start the virtual machine:

   ```
   bin/systemctlstart afo-service.service
   ```

Restarting the virtual machine

Use the following procedure to restart the virtual machine.

Procedure

1. Use SSH to log into the KVM hypervisor as the root user.

2. Use the following command to restart the virtual machine:

   ```
   /bin/systemctlrestart afo-service.service
   ```

Checking the status of the virtual machine

Use the following procedure to check the status of the virtual machine.

Procedure

1. Use SSH to log into the KVM hypervisor as the root user.
2. Use the following command to check the status of the virtual machine:

   ```
   /bin/systemctlstatus afo-service.service
   ```

Chapter 7: User interface troubleshooting

Use the following information to troubleshoot the user interface.

Resetting the admin password on the appliance

Use the following procedure to reset the admin password.

Procedure

1. Go to the following URL: `https://<platform-FQDN>/local-login`
2. From the resulting screen, login as a local administrator with a local administrator password.

   The local administrator means the SSH user account and the corresponding password.

   Only local OS accounts that have been previously designated for emergency login are allowed. You cannot use the network authenticated User IDs.
3. In the resulting screen, enter the User ID and new password.

Unable to login

Condition

Unable to log in to the interface.

Cause

The admin password must be reset.

Solution

1. Go to the following browser URL: `https://<platform-FQDN>/local-login`
2. Log in as a local administrator with a local administrator password.

   The local administrator means the SSH user account and the corresponding password.

Only local OS accounts that were previously designated for emergency login are allowed. You cannot use network authenticated user IDs.

3. In the Password Reset screen, enter the User ID and new password.
4. Confirm the new password.

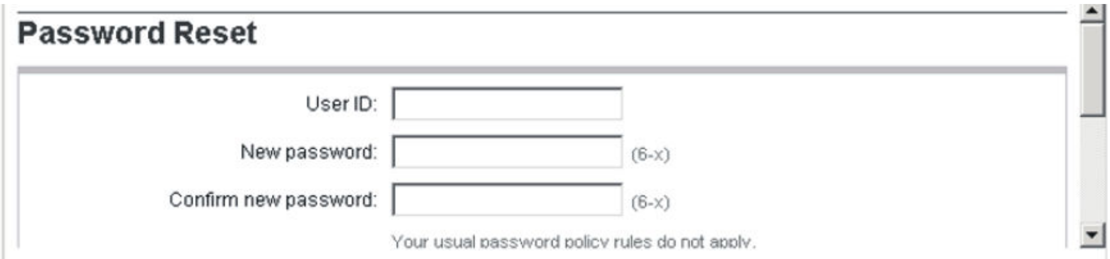

Password Reset

User ID:

New password: (6-x)

Confirm new password: (6-x)

Your usual password policy rules do not apply.

5. Click **Save**.

Condition

Unable to log in to the interface.

Cause

A message indicates the user interface is unavailable.

Solution

Wait for a few minutes and try again.

Condition

Unable to log in to the interface.

Cause

A message indicates an invalid username and password.

Solution

1. Ensure you have entered the proper credentials.
2. Check the cap locks and number locks status on your keyboard.

Unable to login to the user interface

Condition

Unable to log in to the user interface, and you see the following message: `Access Denied: You don’t have permission to access AFO`.

Cause

You may not have the proper permissions.

Solution

1. Check the permissions that you have on the system.
2. Go to the web browser user interface.
3. Select **Administration** > **Users**.
4. Check under the Roles column what role you are assigned:
    - System Administrator—The System Administrator role gives you read-write access across the system, along with modify access to SMGR.
    - Network Administrator—The Network Administrator role gives you read-write access across the system.
    - Network Operator—The Network Operator role gives you read-only access across the system.

Pages do not open

Condition

The selected interface pages do not open and the system displays messages such as the following:

- This page cannot be displayed.
- Unable to connect.
- The connection is untrusted.
- Your connection is not secure.
- There is a problem with the security certificate of this website.
- Content is blocked because the content does not have a valid security certificate.

Cause

The system has its own default certificate authority (CA). Add this CA to the list of trusted CAs for your browser to avoid certificate errors.

Solution

1. Click on **Continue to this website** if the first page that you tried to access was blocked.
2. Select the **About** icon from the quick access toolbar.
3. Click **Install EFO Certificates**.
4. On the About page, click **Download CA Root Certificate** to add the CA to the trusted list of CA, and follow the instructions on the dialog box.

Tabs going missing in tab scope

Condition

Some tabs listed in the tab scope go missing after you open other tabs.

Cause

The tab scope is the drawer for the latest 12 tabs that you have opened, including the home tab. After you open the 13th tab, the system closes the oldest tab on the list. Only 12 tabs can stay in active memory at a time. Items in the tab scope indicate that the tab is active in the background.

After you close an item in the tab scope, the tab closes, and the system halts the activities associated with that particular tab.

The text within the icon displays the number of tabs opened.

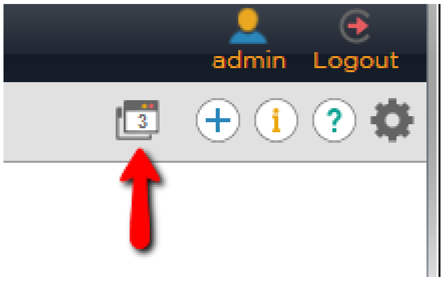

Screen is not rendered properly

Condition

Screen elements appear small without scroll bars and you are not able to view the entire screen.

Solution

1. Set the screen resolution to at least 1024x768.
2. Make sure you are using a supported browser.

Supported browsers and applications

Before deploying the system, ensure that you have the following supported browsers.

Supported browsers

The following section lists the supported browsers.

- Internet Explorer, versions 11.x
- Mozilla Firefox, versions 54 and later

- Apple Safari, macOS 10.8 and later

Supported applications

- Base Operating System — RHEL 7.1, 64-bit
- Hypervisor — Redhat KVM version 7.1
- Virtual Network — Linux bridge/Fabric enabled OpenvSwtich bridge

Client is slow

Condition

Client is slow.

Cause

The browser is consuming too much memory.

Solution

1. Open the **Task Manager** for your computer console, and check how much memory the browser is consuming.
2. If the memory the browser is consuming over 1 GB of memory, consider closing other browser tabs or restart the browser.

Cause

If the browser does not have memory issues, the client machine might not have enough CPU power or memory to run the application.

Solution

Try a different machine to rule out the client hardware.

FAQs

Q:

What are the benefits of a Extreme Fabric Orchestrator (EFO) hardware appliance?

A:

A hardware appliance reduces the number of SKUs that you need to order for an integrated management solution and eliminates the need to maintain and configure your servers.

Extreme Fabric Orchestrator (EFO) is a virtualized solution inside a pre-configured hardware that makes it easier to patch and upgrade.

Q:

How do customers install the Extreme Fabric Orchestrator (EFO) appliance on their network?

A:

First-boot deployment scripts will ask you a series of questions related to the networking as well as the application, and will auto-configure the appliance and virtual machines. For more information, see *Deploying Extreme Fabric Orchestrator*, NN48100–101.

Q:

How to upgrade for customers who have already purchased maintenance contracts for COM, VPS, or VPFM?

A:

If you are an existing customer with existing maintenance contracts for COM, VPS, or VPFM, you will be able to purchase CF Controller appliance (without add-ons) for a discounted price. For more information on data migration from existing application to EFO, see *Deploying Extreme Fabric Orchestrator*, NN48100–101.

Q:

How will customers be able to upgrade EFO in the future?

A:

The EFO appliance includes a Management Server Console application that will help monitor the state of the various virtual machines on the appliance and help to upgrade the virtual machines.

Q:

Can customers deploy their virtual machines on the EFO appliance?

A:

No. You cannot deploy any other virtual machines other than EFO and add-ons on the appliance, as EFO is a closed system.

Q:

Will customers be able to manage the CF Controller appliance using vCenter?

A:

The appliance will not include vCenter. Although you can use your existing vCenter to manage the appliance. As noted above, the appliance is a closed system, and the management virtual machines are not allowed to migrate out of the appliance, and neither any other virtual machines be able to migrate into this appliance.

Chapter 8: Troubleshooting preferences

Use the following information to troubleshoot preferences.

Email notification troubleshooting

Condition

You perform the test email option under **Preference** > **Configuration** > **General**, and receive a status message like the following: `Failed to send Test Email. Error: <error message>.`

Cause

Incorrect information may exist in the **General** tab in the Preferences section.

Solution

1. Go to the web browser user interface.
2. Select the **Preferences** icon from the quick access tool bar to open the Preferences page.

   The Preferences page displays global preferences as the default view on the right side of the page.
3. Select the **Global** tab.
4. Review the primary and backup SMTP hosts configured. The SMTP hosts must be reachable from the configuration virtual machine (VM). Make sure that the SMTP host FQDN is provided in this section. Also make sure that the SMTP IP to FQDN mapping is available in the configuration VM hosts file.
5. Review the SMTP user and password information. Make sure that the SMTP user and password are valid.
6. Review the from and to email users information. Ensure that you provide valid from and to email users information.

Chapter 9: Troubleshooting applications

Use this information to troubleshoot applications on the system.

Applications troubleshooting

Use the following procedures to troubleshoot applications.

One or more applications not responding

Solution

1. Use SSH to log in to the KVM hypervisor as the root user.
2. Query the status of the services using the health check tool:

   ```
   cluster-health-check
   ```

3. For every service that is down, restart the service:

   ```
   cluster-service start -serviceid <platform | fault | flow | config
   | config1 | config2 | msc>
   ```

4. If the services are still not responding, use SSH and log in to the hypervisor again.
5. Run the following command:

   ```
   /bin/systemctl restart afo-service.service
   ```

Applications slow to respond

Condition

The applications on the system are slow to respond.

Solution

1. Use SSH to log in to the KVM hypervisor as the root user.
2. Run the following command to determine the CPU and memory usage of each of the services:

   ```
   cluster-resource-usage
   ```

3. Query the service status using the health check tool:

   ```
   cluster-health-check
   ```

4. When the services are up, check the application response.
5. If the application response has not improved, contact customer support.

Chapter 10: Appliance Device Manager troubleshooting

Use the following information to troubleshoot Appliance Device Manager.

Appliance Device Manager

You can access Appliance Device Manager under **Administration** > **Appliance Device Manager**.

Appliance Device Manager (ADM) is available when the Management Server Console (MSC) server is active. ADM is available using the local login when Single sign-on (SSO) is down.

iLO launch point

ADM allows you to launch Integrated Lights-Out (iLO) from ADM. ILO is a remote management processor that is part of the HP ProLiant server and offers you a virtual presence to access the server from remote sites.

Key Health Indicators

ADM allows you to check the overall health status of the appliance. Three sections exist to check the overall health of the appliance:

- Device Physical View—After you access ADM, the first screen displays the physical device view along with the overall health status of the appliance. The top panel displays a real-time physical view of the front or back panel of the appliance. You can use the **Flip** icon at the top-left corner on the toolbar to rotate the view as front or back. The ADM physical device view indicates the status of the LEDs and the physical components of the KVM server, disks, power supply, and interfaces.

  The module LEDs and the ports are color-coded to provide status. Green indicates the module or port is up and running, red indicates the module or port is disabled, and amber indicates an enabled port that is not connected to anything.

- Services—The services section displays each of the services (virtual machines) and the associated properties of the services on the appliance.
- Monitoring graphs—The monitoring graphs display the statistical view of the CPU usage and memory usage of the services on the device.

Table 1: Monitoring Graphs

Name	Description
Node Status	Displays the statistical view of the hosted data on the server.
Service CPU Usage	Displays the CPU usage of the services on the device.
Service Memory Usage	Displays the memory usage of the services on the device.

Cluster section

The Cluster section in the **Appliance Device Manager** allows you to launch a separate tab that provides key health indicators for the master (secondary server).

Resetting the appliance

Use the following procedure to reset (reboot) the appliance.

✱ **Note:**

Use the **Appliance Reset** button carefully. After you click the button, the whole server reboots, which takes between 10 to 15 minutes.

Procedure

1. Go to the web browser user interface.
2. Select **Administration** > **Appliance Device Manager**.
3. Click the **Appliance Reset** button.

Integrated Lights Out troubleshooting

Use the following information to troubleshoot the Integrated Lights Out feature.

ILO IP address is missing in leader node

Condition

The Integrated Lights Out (iLO) IP address is missing in the leader node (primary node).

Cause

As you load the Appliance Device Manager leader node screen, the following error message appears: `Discovery Failed! Please configure iLO IP address in MSC Preferences`.

If you see the Discovery Failed message, the system is letting you know that the iLO IP address configuration is missing.

Solution

1. Go to the web browser user interface.
2. Click the **Preferences** icon on the quick access toolbar on the top right.
3. On the Preferences page, click **MSC** from the left navigation pane to open the MSC Preferences page.
4. On the MSC Preferences page, click the **iLO** tab.
5. Enter the iLO leader IP address in the **iLO IP** field.
6. Enter the required information into the SNMP v3 section.
7. Click **Apply**.

ILO IP address is missing in master node

Condition

The Integrated Lights Out (iLO) IP address is missing in the master node (secondary node).

Cause

As you load the Appliance Device Manager master node screen, the following error message appears: `Discovery Failed! Please configure iLO IP address in MSC Preferences`.

If you see the Discovery Failed message, the system is letting you know that the iLO IP address configuration is missing.

Solution

1. Go to the web browser user interface.
2. Open a new browser tab using the URL: `https://`*`<master_node_FQDN>`*`/ssd`.

   Note:

   Do not use the IP address.
3. Select **Administration** > **Software Director**.
4. On the Software Director page, click the **MSC Preferences** icon from the top left toolbar.
5. On the Preferences page, click **MSC** from the left navigation pane to open the MSC Preferences page.
6. On the MSC Preferences page, click the **iLO** tab.

7. Enter the iLO master IP address in the **iLO IP** field.
8. Enter the required information into the SNMP v3 section.
9. Click **Apply**.

iLO IP address is not reachable

Condition

After you launch Appliance Device Manager (ADM), you see the following message: `iLO IP Address <A.B.C.D> is not reachable. Functionality dependent on iLO will not work correctly.`

Cause

The message notifies you that iLO IP is not properly configured.

Solution

1. Go to the web browser user interface.
2. Click the **Preferences** icon on the quick access toolbar on the top right.
3. On the Preferences page, click **MSC** from the left navigation pane to open the MSC Preferences page.
4. On the MSC Preferences page, click the **iLO** tab.
5. Enter the iLO leader IP address in the **iLO IP** field.
6. Enter the required information into the SNMP v3 section.
7. Click **Apply**.

Chapter 11: Solution Software Director troubleshooting

Use the following procedures to troubleshoot the Solution Software Director.

Software Director login modes

SSO-based login

When you log in to the platform, you use the single sign-on (SSO)-based login to log on to the Software Director (SSD) as well as the other applications.

Checking the upgrade progress

Use the following procedure to check the upgrade progress.

About this task

If you see a message about session expiry during the upgrade, and the system redirects you to the login page, the upgrade still continues. Upgrade procedures continue until the upgrade completes or fails.

Procedure

1. Log into the web browser user interface.
2. Select **Administration** > **Solution Software Director**.
3. Click **Perform Upgrade in Advanced** mode.
4. Under the Activity Logs section, at the bottom of the page, you can check the progress of the upgrade.

Progress notification lost during upgrade

Condition

If the browser session times out during an upgrade, and the user selects an application, the session is redirected to the login page. The monitoring upgrade progress notification is lost.

When a user logs in again and launches Solution Software Director, the landing page System Status section displays the message: `System upgrade is in progress` and both Easy Mode and Advanced Mode links are disabled. Upgrade activities continue until the upgrade completes or fails.

Cause

During a system upgrade, if the browser session is idle for an extended period of time, the session times out. The system redirects the user to the login page. Only the user session that triggered the upgrade is notified of progress.

Solution

1. Log in to the interface and select **Administration** > **Solution Software Director**.
2. In the Activity Logs section, follow the logs as the upgrade progresses.
3. In the Activity Logs section, click **Save activity logs** to regularly download the activity logs and check for progress.

MSC login redirects to platform login

Condition

You are using the network login to access the dashboard and you want to switch to the MSC login.

Solution

1. Log out of the system.
2. Clear the browser cache.
3. Launch `https://<msc-server-FQDN>/msc-login`

Software Director redirect to MSC

Condition

The system redirects you to the MSC login during the upgrade process in Software Director, and Software Director launches Appliance Device Manager.

Solution

1. Click the Software Director link in the top-right corner of the page to go to Software Director.

2. From the Software Director, you can continue the upgrade procedure.

Unable to upload files to software library

Condition

Errors occur when files are uploaded to the software library.

Cause

File upload failed due to the wrong download ID.

Solution

1. Log in to https://plds.avaya.com/ with your login credentials.
2. Check the download ID of the to-be-uploaded file against the download ID of the corresponding file in PLDS.

Cause

File upload failed due to check-sum mismatch.

Solution

1. Log in to https://plds.avaya.com/ with your login credentials.
2. Check the check-sum of the file to-be-uploaded by running `sha512sum` on the file and confirm that the check-sum matches the check-sum listed on the PLDS entry for the matching file. You need to use your PLDS login credentials.
3. If the check-sum does not match, re-download the file and check the check-sum again.
4. If the file upload continues to fail, contact customer support.

Cause

File upload failed due to wrong file name.

Solution

1. Log in to https://plds.avaya.com/ with your login credentials.
2. Check the expected file name that corresponds to the download ID in PLDS.

Unable to perform prechecks in Solution Software Director

Condition

The system outputs the error `Failed to perform prechecks` in the activity logs in the Solution Software Director.

Solution

Re-run the following script:

```
/opt/avaya/afo/ssd/ha_jms_conf/haSwitchOver.py <Hypervisor Integration IP>
```

For instance, enter something like the following:

```
opt/avaya/afo/ssd/ha_jms_conf/haSwitchOver.py 198.51.100.0
opt/avaya/afo/ssd/ha_jms_conf/haSwitchOver.py 198.51.100.0
```

Chapter 12: Troubleshooting network configuration

Use the information in the following section to troubleshoot network configuration.

Network configuration

To keep track of the network configuration, gather the information described in the following sections. The network configuration information, when kept up-to-date, is extremely helpful for locating information if you experience network or device problems.

Site network map

A site network map identifies where each device is physically located on site, which helps locate the users and applications that a problem affects. You can use the map to systematically search each part of the network for problems.

Logical connections

Ensure that you know how the devices connect logically as well as physically.

Device configuration information

Maintain online and paper copies of the device configuration nformation. Store all online data with the regular data backup for the site.

Other important data about the network

- All passwords—Store passwords in a safe place. A good practice is to keep records of previous passwords in case you must restore a device to a previous software version and need to use the old password that was valid for that version.
- Device inventory—Maintain a device inventory, which lists all devices and relevant information for the network. The inventory allows you to easily see the device type, IP address, ports, MAC addresses, and attached devices.
- Change control—Maintain a change control system for all critical systems. Permanently store change control records.
- Contact details—Store the details of all support contacts, engineer details, and telephone numbers.

Connectivity problems

Use the following general tasks to isolate connectivity problems.

- Check physical connectivity.
- Use tools like ping or trace to verify if the connectivity issue is with an individual port or VLAN.
- Ensure that the URL is in the proper format.
- Try to localize the affected range of ports. If you contact technical support staff to help troubleshoot connectivity problems, always provide source and destination IP address pairs to facilitate troubleshooting. Be sure to provide both working and non-working IP address pairs for comparison.

Host login default credentials

Refer to the following table for a list of default credentials on the solution. A few things to note:

1. Credentials are case sensitive.
2. When you first login to the web browser user interface, the username is admin and the password is admin123.
3. The default root password is Avaya_123. Root access is not allowed to any of the VMs. Switch to root after logging in as admin.

Important:

After first login, change the default passwords to ensure security.

Host	SSH credentials	Local login web credentials
Host (server)	root/Avaya_123	Not applicable
Platform VM	admin/Afo_123	admin/Afo_123
Monitoring VM	admin/Afo_123	Not applicable
Configuration VMs	admin/Afo_123	Not applicable
Flow VM	admin/Afo_123	Not applicable
MSC VM	admin/Afo_123	admin/Afo_123

Unable to launch Configuration pages

Condition

Unable to launch the Configuration pages. Contact your administrator if you see a page with the following message:

```
This request cannot be processed.

No free server instance is available to process this request. Please try
after sometime.

If the problem persists, please logout of the application, login back
and try again.

If the issue still persists, please contact your administrator.
```

Cause

The Configuration pages are not launching because no configuration server is available to process the request. The application servers may have gone down.

Solution

Restart the dashboard JBOSS, and then the remaining CONFIG instances (CONFIG1, CONFIG2, CONFIG3) in order. Select **Administration** > **Appliance Device Manager**. In the Services window, you can select Restart to start JBOSS and the CONFIG instances.

No data displaying in the Network Map tab

Condition

No data is displaying in the **Configuration** > **Network Map** tab.

Solution

1. Verify that the devices have been discovered under **Network** > **Discovery**.
2. Under Discovery, check that there is a positive count for switch (L2), switch (L3), and router.
3. If the count is positive, use the reload map icon in the **Configuration** > **Network Map** view.

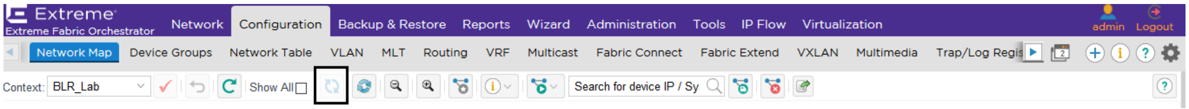

4. If the count is zero, verify the device credentials under **Administration** > **Credentials**.
5. Trigger a new discovery in **Network** > **Discovery**.

Finding the base URL for Configuration Views

Use this procedure to find the base URL for Configuration Views.

Procedure

1. Select **Configuration** > **Network Table**.

2. Select a device, and expand.
3. Move the cursor to the **Links** tab.

   The base URL for **Configuration Views** displays on the status bar of the browser.
4. You can also right-click on any configuration view/panel to view the Frame Info, which also displays the URL.

Troubleshooting device timeouts

Condition

You see SNMP timeout errors under **Configuration** > **Views** after a VLAN or MLT is discovered.

Solution

Ping the device from the COM server. If ping works, then the device is reachable.

Condition

You see SNMP timeout errors under **Configuration** > **Views** after a VLAN or MLT is discovered.

Cause

The SNMP credentials are wrong or SNMP is not enabled on the device.

Solution

1. Right-click the device and select **Show Properties**.
2. If you are able to see **Show Properties**, then you know the device can talk to SNMP. If this fails, then the SNMP credentials are wrong, or SNMP is not enabled on the device.

Condition

You see SNMP timeout errors under **Configuration** > **Views** after a VLAN or MLT is discovered.

Cause

Device is slow to respond or the network latency is high.

Solution

1. To check this, go to the `/opt/avaya/smgr/com/log/COM_Server.log file`, and look for the response after timeout. A value of <- 192.0.2.1, means the device 192.0.2.1 sent the response but the response was received after the SNMP timeout. Increase the SNMP timeout to fix this issue. Go to **Preferences**, and select **Global**. Under the SNMP section, you can update the timeout value in the Timeout field.
2. If the response after timeout message was not present, then Configuration did not receive the message. In this case, use WireShark to capture traffic on the device side. Using WireShark, you can determine if the device did not respond, or if a router in the network

dropped the packet. Routers sometimes drop large packets. In which case, you may want to look at the configuration of the network.

Network Map is slow to render

Condition

The **Network Map** tab under **Configuration** > **Network Map** renders slowly.

Configuration uses mxGraph, a third-party Javascript, to render topology If you have a large network of up to 5000 devices, the system may be slow to render a high number of devices.

The browser may display a popup window asking you to stop running Javascript because Javascript is slow. Do not stop Javascript. Always allow the system to complete the rendering.

Cause

The system may be slow to render a high number of devices.

Solution

1. In Firefox, to eliminate the popup window asking you to stop running Javascript, you can select **Do not ask me this question**.
2. In Internet Explorer, to eliminate the popup box asking you to stop running Javascript, use a registry editor such as Regedt32.exe, and open the following:

   ```
   -key:HKEY_CURRENT_USER\Software\Microsoft\Internet Explorer\Styles
   ```

   Note:

   If the Styles key is not present, create a new key called Styles.
3. Create a new DWORD value called “MaxScriptStatements”, under the Styles key, and configure the value to the required number of script statements. If you are unsure of what value you need to configure, you can configure the value to a DWORD value of 0xFFFFFFFF to avoid this dialog.

   The key does not exist by default. If you do not add the key, the default threshold limit for the timeout dialog box is 5,000,000 for Internet Explorer 4 and later.

Chapter 13: Troubleshooting Bulk Provisioning

The following sections provide troubleshooting information for Bulk Provisioning.

Firewall configuration

Bulk Provisioning uses Telnet, SSH, FTP, SCP, TFTP and SFTP protocols to communicate with various devices and transfer files. If there is a firewall between your devices and the Bulk Provisioning server, you must open up the affected protocols in your firewall configuration.

FTP servers

Do not install FTP servers on a machine on which Bulk Provisioning is installed. Bulk Provisioning starts its own FTP server and installing another FTP server causes the Bulk Provisioning to malfunction. If you experience problems with Bulk Provisioning, uninstall any FTP servers and reboot your machine.

NAT

If you use Network Address Translation (NAT) on your network, ensure that the devices can reach the Configuration server IP address.

Saving CLI correspondence with a device to a file

Procedure

1. Create a new traffic.control file in the COM home folder (`/opt/avaya/smgr/com/`).

> **Tip:**
>
> The traffic.control file is not a text or .txt file.

2. Open the file in text editor.
3. You can edit the file to record traffic for all devices or for selected devices.
   - Option 1: To record traffic for all devices, type ALL on the first line of the traffic.control file and then Save and Close. Files of the form xx.xx.xx.xx.traffic are created in the `/opt/avaya/smgr/com/` folder.
   - Option 2: To record traffic for selected devices, type the IP address of each device on a separate line, and then Save and Close the file.
4. To disable traffic recording, you can delete the traffic.control file or type NONE on the first line of the traffic.control file so you can keep the information in the file.

Terminal length

If you see an unexpected failure of Bulk Provisioning operation with the message “Error while getting device current running image”, then check the terminal length on the device using CLI. If the terminal length is 0, then set the terminal length to a nonzero value. The typical nonzero value is 23.

Configuration e-mail settings

During e-mail configuration, when the **Test Email** button is clicked you may receive an error message stating your anti-virus software is blocking mass e-mail or e-mail worms. This can happen when anti-virus software installed on the Configuration Server is configured to block mass mailing. In order to avoid this, disable the blocking option through the anti-virus software installed on the Configuration server.

Chapter 14: Troubleshooting Virtualization

Use the following information to troubleshoot Virtualization.

Devices missing in inventory because CDP not enabled

Condition

Certain ESX servers are missing from the Inventory.

Cause

CDP is not enabled on the vSwitches or dvSwitch on this ESX server.

Solution

1. To resolve the issue, check to see that CDP is enabled on the vSwitches and dvSwitch on this ESX server.
2. If CDP is not enabled, log in to the ESX server.
3. To list all the vSwtiches on the ESX server, use the following command:

   ```
   esxcfg–info
   ```

4. To enable both mode for CDP on each vSwitch, use the following command:

   ```
   esxcfg-vswitch -b <vSwitch name>
   ```

   This command ensures that all vSwitches are in the BOTH mode to both listen and advertise.

   After CDP is enabled, ESX servers send out CDP packets to the devices, thus being registered in the device forwarding databases (FDBs).

5. To enable CDP on dvSwitch, connect to the vCenter server using the vSphere Client:
   1. In the vCenter server home page, and click **Networking**.
   2. Right-click the **vDS**, and click **Edit Settings**.
   3. Select **Advanced** under Properties.
   4. Select the checkbox to enable CDP, and use the drop-down menu to select both mode.

Note:

LLDP and CDP are supported in dvSwitch

Event category of VM creation and clones results in monitor failures

Condition

Event category of VM creation and clones results in monitor failures.

Causes

- No exact rule matching exists for the newly created VM.
- The rule has criteria for VM_TYPE that cannot be configured until after the VM is created.
- The VM is created on a ESX server and cluster that was recently created, and therefore this ESX server and cluster is not known to the Virtualization inventory.

Solution

1. If the ESX server (or a cluster) is added, Virtualization topology does not automatically learn about the ESX server (or a cluster). Run an Inventory Audit, and (platform) hypervisor connectivity.
2. Most dynamic changes are reported from vCenter through Reconfigure events like PortGroup changes, MAC address changes, ESX server, and cluster name changes. However, if you add new servers, or create an entire new cluster the system does not consider the change as a reconfigure by VMware, hence Virtualization needs to manually rediscover this change.
3. For cases where a VM_TYPE rule is applicable, the VM_TYPE can only be created after the virtual machine (VM) is created, and hence rules with this criteria can fail to execute for the create event.

Event category of VM migration results in monitor failures

Condition

Event category of virtual machine (VM) migration results in monitor failures.

Cause

- No exact rule matching exists for the migrating virtual machine (VM). (This can also be observed during the creation of the VM).
- The VM is migrating to a new ESX server and cluster that was recently created, and therefore this ESX server and cluster is not known to the Virtualization inventory.
- This is a live VM, (powered on).

Solution

1. If the ESX server (or a cluster) is added, Virtualization topology does not automatically learn about the ESX server (or a cluster). Run an Inventory Audit, and (platform) hypervisor connectivity.
2. If this is a live VM (or powered on), vMotion may not be turned on for this ESX server. Live migration does not occur without vMotion.
3. If this is a live VM, (or powered on), the storage used for this VM may not be centrally allocated and attached to the server, hence not allowed to migrate live.
4. If this VM is changing host, and this migration is occurring due to DRS, this migration may not have relevant hosts that share the same PortGroups, and hence migration fails.
5. If vCenter shows the migration as pass, but Virtualization shows failure, look at the following:
    - The **Virtualization** > **Audit Logs** will show the failure reason. Typically this migration monitor will fail, if any of the configurations on the switch have failed. This includes:
        - VLAN creation failure
        - Port association failure
        - Traffic profile creation/ACL failure
    - The Configuration Audit logs will also show what failed on which network switch. If this is a case of the traffic profiles resources being exhausted on the Stackable edge switch, it means more VMs exist on each port than the switch can handle.

Event category of VM delete results in monitor failures

Condition

Event category of virtual machine (VM) delete results in monitor failures.

Cause

- No exact rule matching exists for the deleted VM.
- The VM being deleted has conflicting configurations on the network switch.

Solution

If vCenter shows the deletion as pass, but Virtualization shows failure, following should be looked at:

- The Virtualization Audit logs will show the failure reason. Typically this migration monitor will fail, if any of the configurations on the switch have failed. This includes:
    - VLAN delete failure
    - Port disassociation failure
    - Tarffic profile delete/ACL failure

The Configuration Audit logs will also show what failed on which network switch. Typically when the same VLAN, port, or Traffic profile is in use by other virtual machines (VMs), these will not be deleted from the network switch, and not be marked as failure.

If other VMs in the network use the same VLAN, the VLAN is not deleted from the network switch. The same is the case for port removal from the VLAN. The port is only removed if other VMs are not using the port.

Event category of VM reconfigure results in monitor failures

Condition

Event category of virtual machine (VM) reconfigure results in monitor failures

Cause

The system receives a flurry of events because some reconfiguration on an EXS server causes several VMs to be affected.

Solution

If vCenter shows the reconfigures as pass, but Configuration shows as failure, consider the following:

- The Configuration audit logs will show the failure reason. Go to **Reports** > **Audit Logs** for configuration audit logs. Typically the reconfiguration monitor fails because the system received too many related events and these bulk changes were lost. Run an Inventory Audit, followed by a Hypervisor Connectivity, to allow Configuration to learn the inventory again.
- Note that this is a very rare scenario and can only happen if the admin is knowingly making some basic changes to the entire cluster and the ESX server.

Link information of ESX server to physical network is not visible in the topology or Inventory View

Condition

Link information of the ESX server to the physical network is not visible in the topology or Inventory View.

Cause

- Certain ESX servers are not connected to the physical switches in the topology view.
- SNMP is not enabled on the ESX server.
- CDP is not enabled on the vSwitch or dvSwitch on the ESX server.

Solution

1. Check if SNMP is enabled on the ESX server:

   ```
   esxcli system snmp get
   ```

2. If SNMP is not enabled, use the following commands to enable SNMP on the ESX server:

   ```
   esxcli system snmp set –e yes
   ```

   ```
   esxcli system snmp set –c public
   ```

3. Check if CDP is enabled on the vSwitch and dvSwitch on the ESX server. To enable CDP on the vSwitch or dvSwitch, do the following:

   a. Log in to the ESX server.

   b. Check the configuration of vSwitch on the ESX server using the following command:

      ```
      esxcfg-vswitch –b <vSwitch name>
      ```

   c. Enable CDP on each vSwitch using the following command:

      ```
      esxcfg-vswitch –B both <vSwitch name>
      ```

      This ensures that all vSwitches are in the BOTH mode to listen and to advertise. fter CDP is enabled, ESX server send out CDP packets to the devices, thus being registered in the device forwarding databases (FDBs).

   d. To enable CDP on dvSwitch, edit the dvSwitch configuration settings from the vCenter client to the following settings under the **Properties** tab:

      a. Select **Advanced**.

      b. Enter the **Maximum MTU** as 1500.

      c. Select **Cisco Discovery Protocol**.

      d. Select **Both**.

Event category of VM Events show pending or failed in monitor events

Condition

Event category of VM Events display as pending or failed in monitor events.

Cause

- No exact rule matching exists for the VM Event for the source server.
- Virtualization does not contain the physical switch information.
- The VM Event is on an ESX server and cluster that was recently created, and therefore this ESX server and cluster is not known to the Virtualization inventory.

Solution

1. Create the required network profile to apply the configuration on the physical network and associate it with a matching rule.

2. If the ESX server or cluster is added, the Virtualization topology does not automatically learn of the ESX server and cluster. Run Inventory Audit, followed by Hypervisor Connectivity.
3. The Virtualization audit logs will show the failure reason. Select **Virtualization** > **Audit Logs** to access the Virtualization audit logs. The Configuration Audit Logs will also show what exactly failed on which network switch. Select **Reports** > **Audit Logs** to access the configuration audit logs.

Chapter 15: Troubleshooting IP Flow

Use the following information to troubleshoot IP Flow.

Device added to IP Flow not active or no traffic information displaying

Condition

A device added to IP Flow is not active or no traffic information is displaying on the user interface.

Solution

1. Check the validity of the device IP address in the flow dashboard.
2. Check that the correct IP address is added in the device as IPFIX exporter, which must be the same as the IP Flow server machine address.
3. Check that the UDP ports are defined in the IP Flow preferences.
4. Make sure that no firewall is active between the device and the system.
5. Verify the credentials to access the device. Go to **Administration** > **Credentials** to check credentials.

User does not receive any email with threshold email setup

Condition

User does not receive any email with threshold email setup.

Solution

1. Check that there is valid SMTP information in Global Preferences.
2. Make sure no firewall is active between the SMTP server and the IP Flow server.
3. Ensure that the threshold definition is properly configured for the related IPFIX device. Go to **IP Administration** > **Administration** > **Thresholds**.
4. Ensure that enough traffic exists in the current network so that the threshold is reached.

User cannot make a packet capture from IP Flow

Condition

A user cannot make a packet capture from IP Flow.

Cause

The packet capture feature is only supported by the ERS 8600 device family but not by other device families.

Solution

1. Perform the following checks on the ERS 8600 device:
    - The device must have Dual CPU (8692 cards).
        - CLI command: **`show sys info`**
        - CLI command: **`show boot config general`** or **`show running-config`**
    - The boot config flag must have ha-cpu configured to false.
        - CLI command: **`show bootconfig flags`**
        - CLI command: **`show boot config flags`**
    - The boot config flag must have ftpd configured to true.
        - CLI command: **`show bootconfig flags`**
        - CLI command: **`show boot config flags`**
    - The device must use file capture mode configured to use PCMCIA device
        - CLI command:
        - CLI command:
    - The device must have PCMCIA cards inserted into each CPU slot.
        - CLI command:
        - CLI command:
    - The PCAP file size must be configured to the minimum value of 2 MB.
        - CLI command:**`show diag pcap info`**
        - CLI command: **`show diag pcap info`**
2. Perform the following checks on the IP Flow server:
    - Check that a little traffic is collected before you invoke PCAP operation.
    - Check that a valid SNMP (v1 or v3) and FTP credentials are added. Select **Administration** > **Credentials** to access credential information.
    - Check that the Packet Capture duration value provided in the IP Flow preferences is long enough for the device to generate a PCAP file.

IP Flow cannot meet a traffic burst

Condition

IP Flow cannot meet the speed, such as a traffic burst.

Solution

1. If there is an extremely high volume of incoming traffic, IP Flow is not able to process incoming data efficiently.
2. You can fine tune the following parameters in the `/opt/avaya/smgr/ipfm/ipfix_collector/etc/IPFMCollectorConfiguration.properties` file.

   Take the following steps in the file:

   - Increase the sampling rate under sample.rate.
   - Reduce the data retention time under collector.data.retention.interval.
   - Reduce the cleanup interval under collector.database.cleanup.interval.

Chapter 16: Troubleshooting high availability

Use the following information to troubleshoot high availability.

High Availability

High Availability (HA) ensures devices recover quickly from a failure by:

- Eliminating single points of failure.
- Detecting failures.

To ensure high availability and redundancy, the system uses two SDN appliances on two HP ProLiant DL360 Gen9 servers.

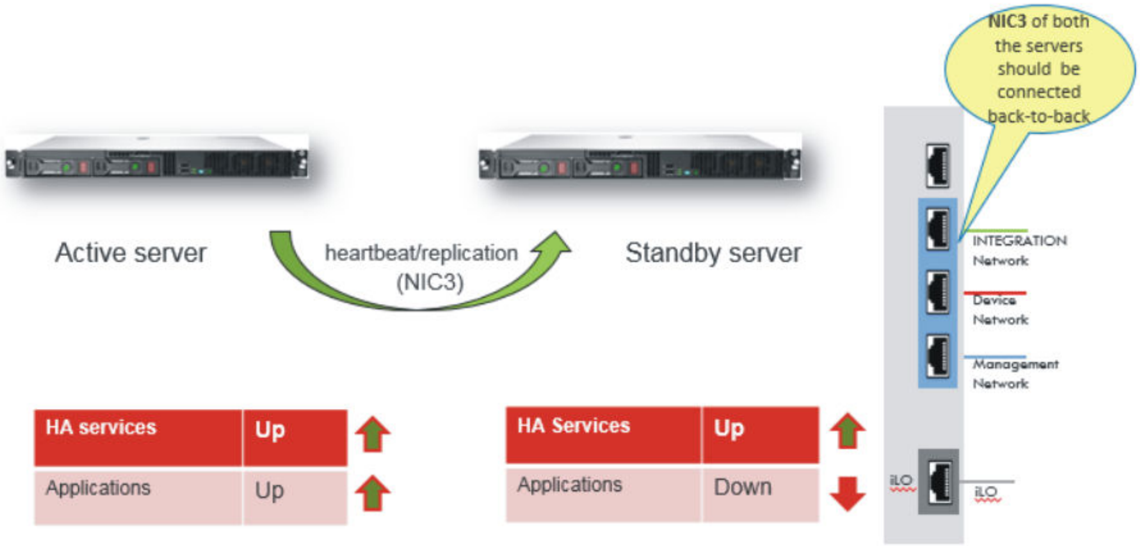

Active-Standby

The virtual machines of the appliance that function in active stand-by are the following: platform, config, and monitoring virtual machines (VMs).

As part of active-standby a fully redundant instance of the appliance and server node takes over if the primary node fails.

In active-standby for the solution, one appliance and server is the primary server (known as the leader in the solution) and is active. A second appliance and server (known as the master in the solution) is inactive.

If the cluster engine determines one of the VMs needs to restart in the case of a failover of one of the system components, the information from the primary server (the leader), is replicated on the secondary server (the master), which then becomes the primary server.

The load balancer also runs in active-active. The load balancer clusters the servers with a single virtual IP address and distributes client requests between those servers dependent on specific criteria. If the load balancer experiences a failover, the secondary load balancer takes over the primary load balancer position.

Active-Active

The Management Server Console (MSC), OpenDayLight Controller (ODLC), and SDN Engine VMs on both the leader (primary) and master (secondary) servers function in active-active. In active-active mode, all three VMs run on both the leader (primary) and master (secondary) appliances and servers.

Cluster engine

The cluster engines on the primary and secondary server act as watchdogs, watching the system to ensure the various services or VMs are running properly. The cluster engines consists of three parts:

- Cluster engine virtual machine
- Cluster engine client
- Cluster engine server

The cluster engines appear as one, and share a management virtual IP address. When you log into the virtual IP address through VNC, you log into the primary cluster engine (the leader master), but you have access to a view of both cluster engines.

The cluster engine does the following:

- Monitors the health of all virtual machines and processes.
- Takes corrective measures if a process or virtual machine fails.
- Applies the corrective mechanism.
- Detects planned and unplanned failovers.
- Performs the failover.
- Provides API interfacing for clients to get/set clustering parameters.

Load balancer

The load balancer is a clustering technology that uses a distributed algorithm to load balance network traffic, enhancing the scalability and availability of mission critical operations.

The load balancer also provides high availability by detecting failures and automatically redistributing traffic to the remaining operational VMs. So, for instance, if one of the SDN Engines fails, the load balancer detects the failure, and ensures the other SDN Engine takes over.

The load balancer itself runs in active-active mode. If the primary load balancer fails, the secondary load balancer takes over and becomes the primary load balancer.

Viewing the High Availability status

Use this procedure to view the High Availability status.

Procedure

1. SSH to the Leader (primary) node or the Master (secondary) node KVM hypervisor as the root user.
2. Run the following command:

   /usr/local/infra/bin/ha_status.sh

Failure or shutdown of the leader node

Condition

In the case of a failure or shutdown of the leader node, the failover from the leader to the master node is automatically triggered.

However, the failback is not automatically triggered. If you want the earlier node to become active again, you must perform a manual failback using the following procedure.

Note:

If you perform the failback procedure, expect a downtime of services for 10 minutes.

Solution

1. SSH to the cluster engine using the following command from the active or standby KVM:

   ```
   ssh admin@192.0.2.1
   ```

2. Enter the following commands in order to perform a forceful failback:
   a. `/cfg/sys/supervisor/failover on`
   b. `apply`
   c. `exit`

Troubleshooting HA split-brain scenario

Condition

Both nodes have become primary nodes, which is also known as a split brain scenario.

Cause

- The integration network is removed or disconnected during data replication.
- The time is not synced between the master (secondary) and leader (primary) nodes, and failover occurs.
- A network failure occurs during a reboot of a node, before services shutdown completely.

Solution

The dashboard shows a HA warning indicating a split-brain scenario. The warning message is repeated every 15 minutes until resolved.

```
Network Split Detected. node40-kvm-master.avaya.com and node40-
kvm.avaya.com are not HA pair anymore.
Please rectify this immediately.
Login to KVM:
    - Execute: bash /usr/local/infra/bin/recover_split_brain.sh
    - Select the Survivor node on prompt
and wait for completion.
```

Run the script as root user, select a survivor node as prompted and confirm your selection. Confirm the prompts to discard and overwrite data on the node to be recovered from the split-brain scenario.

Activating maintenance mode

Use the following procedure to activate maintenance mode. Maintenance mode offers a supervisor a control-free window, during which:

- No heartbeats are sent for health check.
- If a VM or application is offline the system does not indicate a failure.

Activate maintenance mode before performing a system backup, restore, or maintenance.

Procedure

1. Use SSH and the leader IP address to log in to the leader hypervisor as the root user.
2. Use the following command to SSH to the Cluster Engine:

   ```
   ssh root@<Cluster Engine IP address>
   ```

3. Activate maintenance mode in one of two ways on the Cluster Engine:

   a. Use the following command in the CE CLI menu to activate maintenance mode:

      ```
      /cfg/sys/supervisor/maint on
      ```

      OR

   b. Use a RabbitMQ API request: maintenance_req:

      ```
      status=1 to enable maintenance mode
      ```

      ```
      status=2 to check current maintenance mode setting
      ```

4. Deactivate maintenance mode in one of two ways:
    1. Use the following command in the CE CLI menu to deactivate maintenance mode:

       ```
       /cfg/sys/supervisor/maint off
       ```

       OR
    2. Use a RabbitMQ API request: maintenance_req:

       ```
       status=0 to disable maintenance mode
       ```

FAQ — High Availability

What do I do if I see a RedHat Kernal panic message in the platform VM during the boot up, after an HA failover scenario.

- Redhat is designed to be resilient for failure and auto-recovers the system during bootup. Wait for approximately 10 minutes for Redhat to auto-recover the system.

What do I do if the platform VM does not come up after an HA switchover or failover operation?

- The cause can be that the portgres SQL database service on the platform VM does not come up. Check if the postgres is running on the platform VM using the service postgreql status. If the service is not running, check the logs in `var/lib/pgsql/pgstartup.log and /var/lib/pgsql/data/pg_log/postgres.log`

  ```
  Mar 17 02:44:48 cluster-platform postgres[15503]: [2-1] 2016-03-17 02:44:48 EDT:
  [15503]LOG: database system was interrupted; last known up at 2016-03-16 09:49:15
  EDT Mar 17 02:44:48 cluster-platform postgres[15503]: [3-1] 2016-03-17 02:44:48 EDT:
  [15503]LOG: invalid magic number 0000 in log file 0, segment 2, offset 7512064 Mar
  17 02:44:48 cluster-platform postgres[15503]: [4-1] 2016-03-17 02:44:48 EDT:
  [15503]LOG: invalid primary checkpoint record Mar 17 02:44:48 cluster-platform
  postgres[15503]: [5-1] 2016-03-17 02:44:48 EDT:[15503]LOG: invalid secondary
  checkpoint record Mar 17 02:44:48 cluster-platform postgres[15503]: [6-1]
  2016-03-17 02:44:48 EDT:[15503]PANIC: could not locate a valid checkpoint record
  Mar 17 02:44:48 cluster-platform postgres[15494]: [2-1] 2016-03-17 02:44:48 EDT:
  [15494]LOG: startup process (PID 15503) was terminated by signal 6: Aborted Mar 17
  02:44:48 cluster-platform postgres[15494]: [3-1] 2016-03-17 02:44:48 EDT:
  [15494]LOG: aborting startup due to startup process failure
  ```

  If the logs resemble those above, the transaction logs or write ahead logs in postgres SQL are corrupted. Complete the following steps to recover the database.

  1. Login in as the root user to the platform VM.
  2. Enter the following commands:

     ```
     mkdir -p /root/pg_data;cp -r /var/lib/pgsql/data/ /root/pg_data/
     ```

     ```
     udo -H -u postgres bash -c "/usr/pgsql-9.3/bin/
     pg_resetxlog /var/lib/pgsql/data/" service postgresql start
     ```

  3. If the postgres SQL does not come up, contact customer support.

Chapter 17: Log collection

One-click log collection

You can run a single command to collect all application logs (debug/trace/operational/audit/security) into a single archive, for all issues related to troubleshooting. The command collects TFP logs on configuration, domain level information on monitoring, and system information. All logs are archived at a services level, and then at the cluster level.

Use the information in the current section to perform one-click log collection for application logs generated on the centralized server on demand in a single click.

The one-click log collection feature is developed by extending the existing **`createLogArchive.sh`** command which collects log files and other required configurations from individual applications.

The command is executed remotely from the Management Server Console (MSC) on every applicable system and collects the archive to a central place.

One-click log collection configuration

About this task

Use the following procedure to create a log archive of the entire appliance that contains logs of all of its services. The information includes:

- Application specific logs
- Jboss logs
- Avaya CLF logs (operational, audit, and security) files

Logs are hierarchically archived at a service level, and also at an appliance level.

The current procedure collects logs into a single archive from all the applications deployed on the system.

Procedure

1. Use SSH to login to the MSC as an admin user.
2. Run the `su — root` command on the Command Line Interface (CLI) to switch to the root user.

3. Run the following command on MSC server to collect logs:

   ```
   /opt/avaya/afo/infra/OneClickLogCollect.sh
   ```

4. The collected logs are compressed into a single archive and stored in the directory.

   The filename of the generated archive is displayed on the system.

   Note:

   Each time you run the command, the system generates a new zip file at the same location.

   The ZIP file contains the logs from all of the modules, including Monitoring, Configuration, Flow, MSC, and Platform logs.

Chapter 18: Backup and restore troubleshooting

Use the information in the following section to troubleshoot backup and restore issues.

Failover is not allowed error message

Condition

In a failover scenario, when you try to put the Cluster Engine into maintenance mode, you can see the following error message: `Changing Maintenance mode to on while Failover mode is off is not allowed. You must perform a manual failover migration first.`

Cause

Backup and restore functionality is not supported when the cluster is not in a healthy state in failover mode.

Solution

1. On the user interface, select **Administration** > **Appliance Device Manager**.
2. In the Cluster section, enable the failover mode.
3. After the system returns to a cluster healthy state with failover mode enabled, perform a backup and restore.

Pre-validation checks do not continue if one VM is down

Condition

Backup and restore pre-validation does not continue if one of the virtual machines (VMs) is down in the cluster.

Note:

Pre-validation does not check the status of the cluster engine and load balancer.

Cause

Backup and restore functionality is supported only when all of the VMs are in the running state.

Solution

1. Use SSH and the leader hypervisor IP address to log in to the leader hypervisor as the root user.
2. Use the following command to ensure all of the VMs are running:

   ```
   virsh list -all
   ```

   ```
   login as: root
   Access denied
   root@192.0.2.1's password:
   Last login: Fri Jun  3 08:46:21 2017 from 198.51.100.0
   [root@perf-cluster1-kvm ~]# virsh list --all
    Id     Name                           State
   -----------------------------------------------------
    2      MSC                            running
    3      Platform                       running
    4      Monitoring                     running
    5      ODL_Controller                 running
    6      Config                         running
    7      SDN_Engines                    running
    8      Cluster_Engine                 running
    9      LoadBalancer                   running
   ```

3. To restart applications, see Starting an application service on page 29.

Backup and restore does not continue if cluster engine is down

Condition

Backup and restore does not continue if cluster engine is down.

Cause

The backup and restore flow uses the cluster engine APIs (RabbitMQ server) to check which server is the leader (primary) in the cluster. When the cluster engine is down, the flow does not continue.

Solution

1. Check the status of cluster engines in the cluster and make sure the cluster engine is running.
2. Use SSH and the leader hypervisor IP address to log in to the leader hypervisor as the root user.
3. Use the following command to ensure all of the VMs are running:

   ```
   virsh list -all
   ```

   ```
   login as: root
   Access denied
   ```

```
root@192.0.2.1's password:
Last login: Fri Jun  3 08:46:21 2017 from 198.51.100.0
[root@perf-cluster1-kvm ~]# virsh list --all
 Id    Name                          State
----------------------------------------------------
 2     AFO-MSC-1.1.0.0.164                            running
 3     AFO-Platform-1.1.0.0.164                       running
 4     AFO-Monitoring-1.1.0.0.164                     running
 5     AFO-Config-I2-1.1.0.0.164                      running
 6     AFO-IPFLOW-1.1.0.0.164                         running
 7     AFO-Config-I1-1.1.0.0.164                      running
 8     AFO-Config_I3_1.1.0.0.164                      running
 9     AFO-ADS-1.1.0.0.164                            running
10     Cluster_Engine                                 running
```

4. To restart applications, see Starting an application service on page 29.