Universal Compute Platform v5.14.01 User Guide

System Configuration and Management

9041050-00 Rev.AA
June 2026

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, see: https://www.extremenetworks.com/about-extreme-networks/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: https://www.extremenetworks.com/support/policies/open-source-declaration/

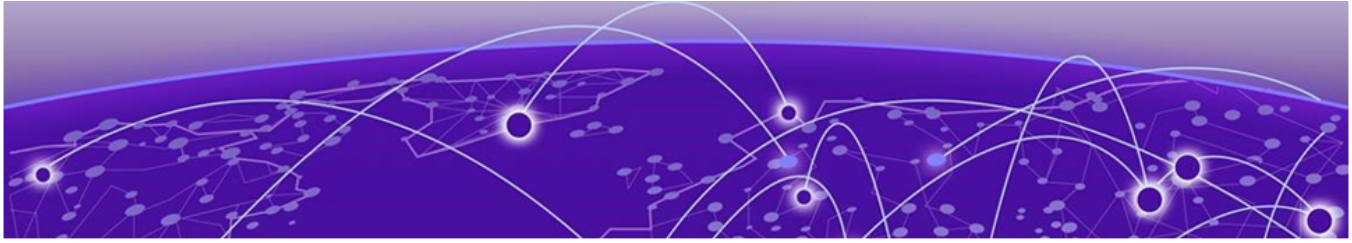

Table of Contents

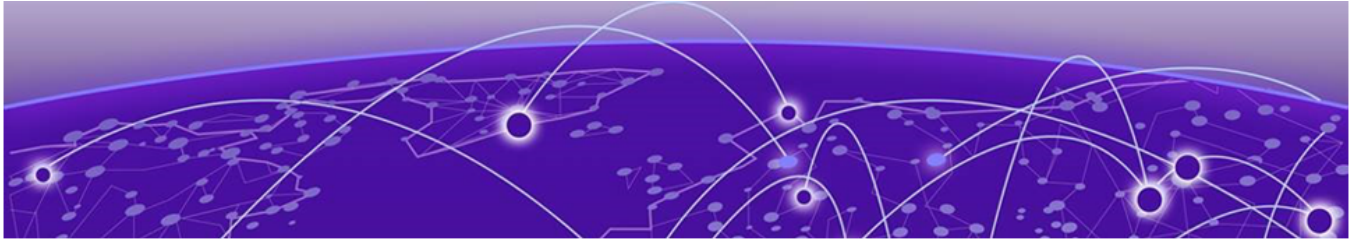

Abstract

This Universal Compute Platform version 5.14.01 User Guide provides in-depth procedures for configuring, managing, and upgrading the Universal Compute Platform environment. The v5.14.01 version of the guide features revisions to the Add Nodes and Replace Nodes procedures, updates to Kubernetes upgrades, support for Access Events logs, simplified VRRP configuration, and the ability to install public CA certificates for selected interfaces. The User Guide details advanced tasks such as user account management, network interface configuration, cluster node setup, and performing system-wide backups. The guide includes comprehensive instructions on engine application configuration, image management, and container orchestration using ExtremeCloud Edge, supporting both managed and self-orchestrated deployments. It also covers troubleshooting protocols, system logging mechanisms, and maintenance strategies with a focus on redundancy configurations, network failover, and high-availability setups. Detailed processes for initiating on-demand backups, scheduling regular backups, and upgrading critical system components are provided to ensure optimal platform performance and continuity, specifically addressing the needs of system administrators overseeing large-scale Universal Compute Platform deployments.

Preface

Read the following topics to learn about:

- The meanings of text formats used in this document.
- Where you can find additional information and help.
- How to reach us with questions and comments.

Text Conventions

Unless otherwise noted, information in this document applies to all supported environments for the products in question. Exceptions, like command keywords associated with a specific software version, are identified in the text.

When a feature, function, or operation pertains to a specific hardware product, the product name is used. When features, functions, and operations are the same across an entire product family, such as Extreme Networks switches, the product is referred to as *the switch*.

Table 1: Notes and warnings

Icon	Notice type	Alerts you to...
	Tip	Helpful tips and notices for using the product
	Note	Useful information or instructions
	Important	Important features or instructions
	Caution	Risk of personal injury, system damage, or loss of data
	Warning	Risk of severe personal injury

Table 2: Text

Convention	Description
`screen displays`	This typeface indicates command syntax, or represents information as it is displayed on the screen.
The words *enter* and *type*	When you see the word *enter* in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says *type*.
Key names	Key names are written in boldface, for example **Ctrl** or **Esc**. If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press **Ctrl**+**Alt**+**Del**
Words in italicized type	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles.
NEW!	New information. In a PDF, this is searchable text.

Table 3: Command syntax

Convention	Description
bold text	Bold text indicates command names, keywords, and command options.
italic text	Italic text indicates variable content.
`[ ]`	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
`{ x | y | z }`	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x | **y**	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, such as passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, *member*[*member*...].
\	In command examples, the backslash indicates a “soft” line break. When a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Documentation and Training

Find Extreme Networks product information at the following locations:

Current Product Documentation
Release Notes
Hardware and Software Compatibility for Extreme Networks products
Extreme Optics Compatibility
Other Resources such as articles, white papers, and case studies

Open Source Declarations

Some software files have been licensed under certain open source licenses. Information is available on the Open Source Declaration page.

Training

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit the Extreme Networks Training page.

Help and Support

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

Extreme Portal

Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

The Hub

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

Call GTAC

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2800. For the support phone number in your country, visit www.extremenetworks.com/support/contact.

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

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions 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

Subscribe to Product Announcements

You can subscribe to email notifications for product and software release announcements, Field Notices, and Vulnerability Notices.

1. Go to The Hub.
2. In the list of categories, expand the **Product Announcements** list.
3. Select a product for which you would like to receive notifications.
4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

You can modify your product selections or unsubscribe at any time.

Send Feedback

The User Enablement team at Extreme Networks has made every effort to ensure that this document is accurate, complete, and easy to use. We strive to improve our documentation to help you in your work, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information.
- Broken links or usability issues.

To send feedback, email us at Product-Documentation@extremenetworks.com.

Provide as much detail as possible including the publication title, topic heading, and page number (if applicable), along with your comments and suggestions for improvement.

Welcome to Universal Compute Platform

Navigating the User Interface on page 10

The Universal Compute Platform serves as a service platform for an on-premises application offering. The Universal Compute Platform provides a performance validated hosting platform, supporting advanced orchestration of a catalog of applications. The Universal Compute Platform provides a container-based orchestration framework, in an Extreme Networks qualified and validated high-performance hardware configuration. The framework natively supports clustering, a distributed file system, and orchestration through Kubernetes, providing a highly resilient application operational base.

Navigating the User Interface

To open the navigation menu for the Universal Compute Platform user interface, select the navigation icon () from the top left of the interface header. From the navigation menu, you can select one of the following menu options to open a page with relevant information.

- Dashboard
- Cluster Settings
  - Cluster Configuration
  - Node Replacement
  - Add Nodes
- Engines
  - Installation
  - Image Management
- Tools
  - Logs
  - Diagnostics
- Administration
  - Accounts
  - System
    - Configuration
    - Logs
    - Maintenance

- Network Setup
- Network Time
- Settings
- Software Upgrade
- System Information
- Utilities

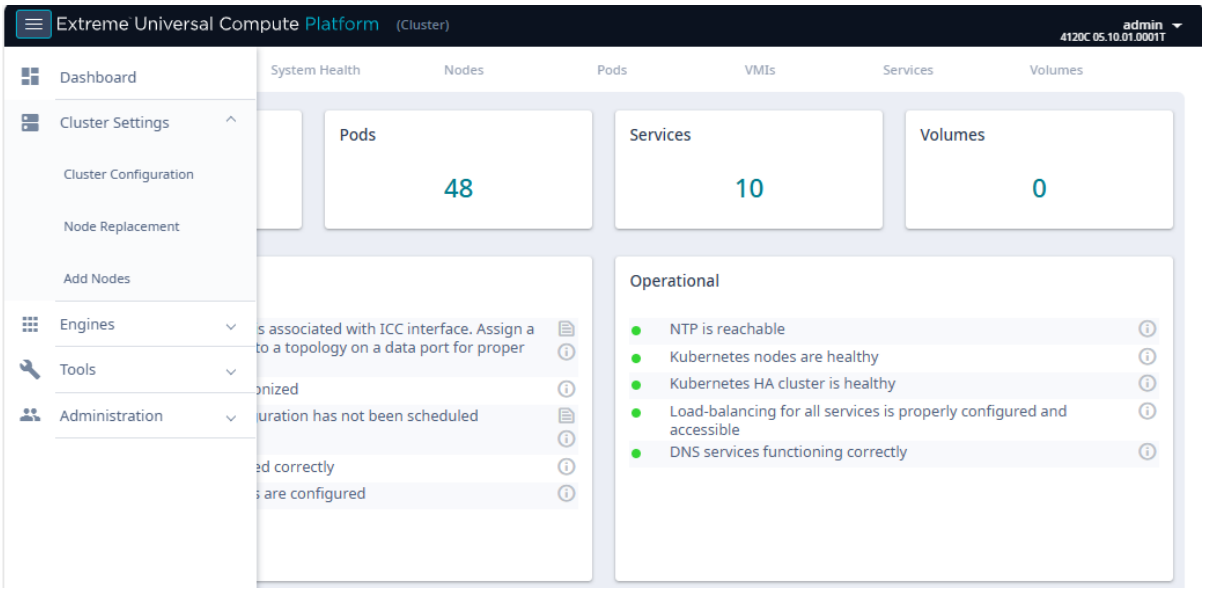

Figure 1: The Universal Compute Platform desktop

Dashboard

The topics in this section describe the dashboards that are available when you select the **Dashboard** menu option.

Dashboard Overview

Universal Compute Platform offers dashboards and lists that help you monitor the cluster configuration and performance.

Universal Compute Platform offers the following dashboards and reports:

- Deployment Health
- System Health
- Dashboard Nodes
- Availability Zones
- Pods List
- Services List
- Volumes List

Deployment Health

The **Deployment Health** Dashboard provides information about the overall health of the node cluster. The top pane highlights each piece of the cluster network:

- Nodes—The number of appliances in your network. You have the option of configuring individual stand-alone nodes or a cluster of three or more nodes. Stand-alone configuration is supported for all engine types except ExtremeCloud IQ.

  Note
  When using an ExtremeCloud™ IQ engine, you must configure a cluster of three or more nodes in multiples of three (for example, three, six, or nine nodes). ExtremeCloud IQ is not supported in stand-alone mode, requires a cluster, and does not support engine types other than ExtremeCloud IQ.

- Pods—A group of managed containers that share networking and storage resources from the same node (appliance). Each pod is assigned an IP address. All the containers in the pod share the same storage, IP address, and network namespace.

- Services—Network Services running on the node cluster.
- Volumes—Storage that allows data to be accessible to containers within a pod.

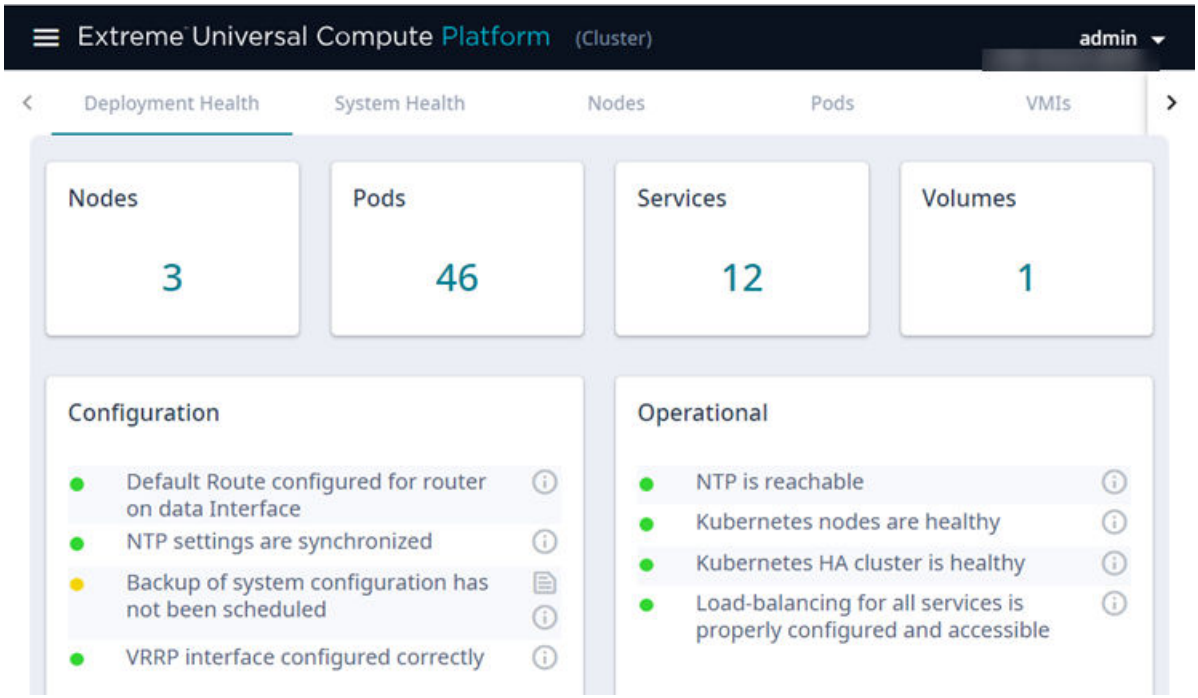

Figure 2: Deployment Health Dashboard

Deployment Health also provides best practice information for your Universal Compute Platform configuration. System Health checks are run against your configuration and operational setup to inform you of best practices.

Examples of what's checked include:

- Default route configuration
- NTP settings are configured.
- System backup is scheduled
- VRRP interface is configured properly.
- Kubernetes nodes and HA cluster are healthy.
- Load balancing is configured.
- DNS settings are configured with Primary and Secondary DNS servers. In addition, DNS is able to resolve URLs successfully.
- Inter-node connectivity (latency) is within acceptable limits.
- All nodes are running the same firmware version.

Health check results are reported using the following scheme:

- Green indicates that a best practice is being followed.
- Yellow indicates that your configuration is not optimal.
- Red indicates an error in your configuration.

Fix all error conditions. You have the option to ignore warnings. They are provided to inform and encourage best practice configuration.

- Select for a description of each statement or warning.
- Select to list objects causing an issue, and to jump to that area of Universal Compute Platform to improve your configuration.

System Health Dashboard

The **System Health** dashboard provides the following information:

- System Uptime — The number of days and hours the system has been operational.
- CPU Utilization — CPU utilization metrics over time. Hover your cursor over the timeline to view the recorded CPU utilization at that interval.
- Memory Utilization — Memory utilization metrics over time. Hover your cursor over the timeline to view the recorded memory utilization at that interval.

Nodes Dashboard

The **Nodes Dashboard** provides graphs for CPU utilization and memory utilization for each node in the cluster. The dashboard also provides the Inter-Node Connectivity Matrix, which reports the per-node results of the latest inter-node connectivity and latency checks.

Availability Zones Dashboard

The **Availability Zones Dashboard** displays a mapping of the Availability Zones that are deployed in the cluster, and which nodes belong to which zone.

The Node column displays the hostname of each node. The node's ICC IP address displays in the column for the Availability Zone to which that node belongs.

Pods List

The **Pods List** displays a list of pods in your cluster. A pod is a group of managed containers that share networking and storage resources from the same node. The following information is provided for each pod:

- Pod Name
- Ready status
- Status — Possible values are Running or Down.
- Restarts
- Age — Measured in minutes, hours, and days.
- IP address
- Node

Use the Search field to find a specific list item.

You can select the number of items to display on a page. Valid values are:

- 10
- 20
- 50
- 100

To jump to the next or previous page, or jump to the first or last page, select the arrows at the bottom of the page.

VMI List

VMI stands for Virtual Machine Instance. The following information is provided for each VMI:

- Name
- Phase
- Node Name
- QoS Class
- Namespace
- Created

Expand each VMI to display the following information:

- CPU: Cores
- Volumes
- Interfaces: IP Address and MAC

Services List

The **Services List** displays a list of all services running in the cluster. The Service Name and Namespace are provided for each service.

Use the Search field to find a specific list item.

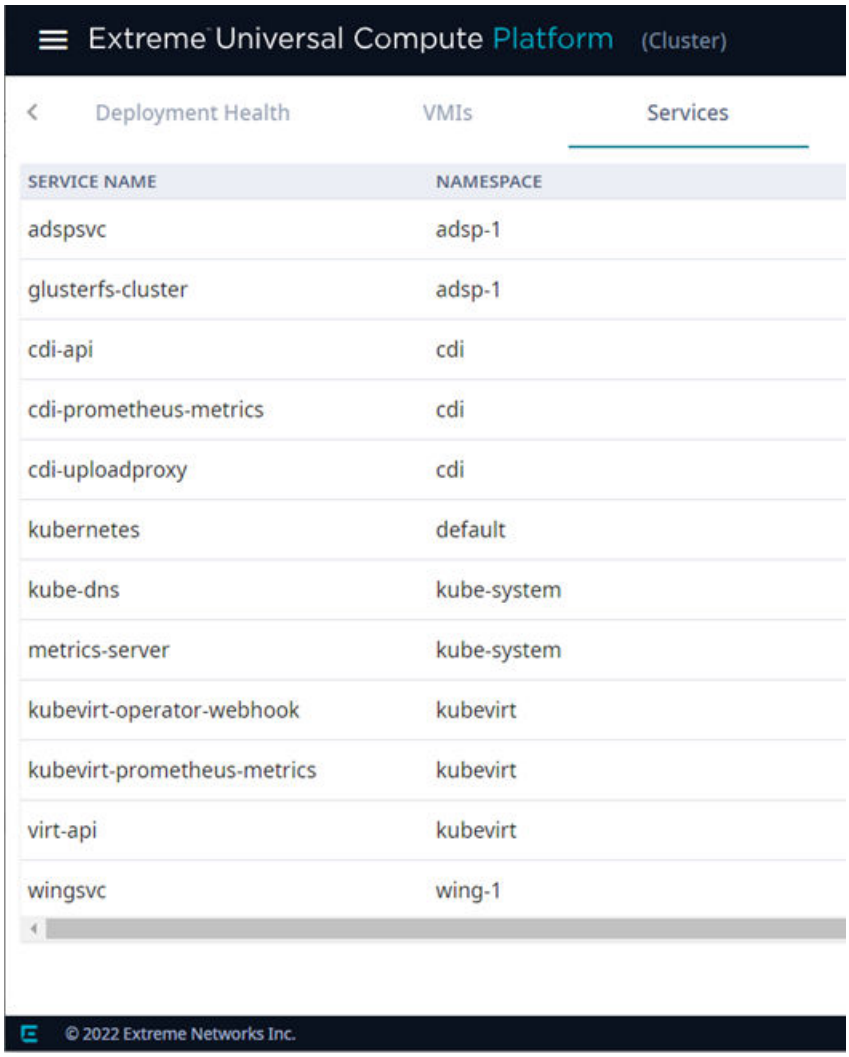

Figure 3: List of services running on the node cluster

You can select the number of items to display on a page. Valid values are:

- 10
- 20
- 50
- 100

To jump to the next or previous page, or jump to the first or last page, select the arrows at the bottom of the page.

Volumes List

The **Volumes List** displays a list of all volumes in the cluster. A volume is storage that allows data to be accessible to containers within a pod. The following information is provided for each volume:

- Volume Name
- Capacity

- Status
- Claim. Associated with the volume type and how the data is handled in the volume. If the data will be retained, the Claim value is **Retain**.

Use the Search field to find a specific list item.

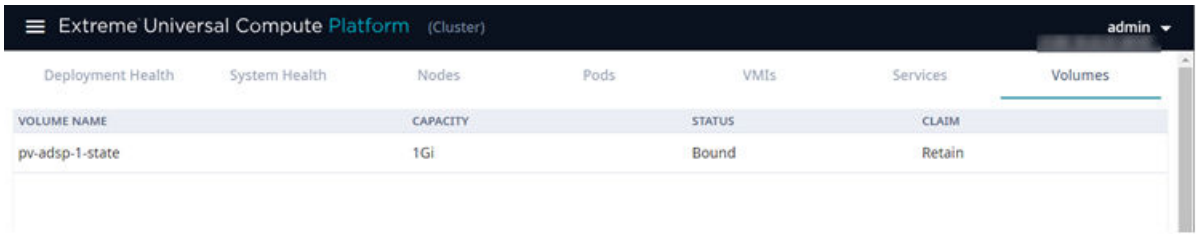

Figure 4: List of Volumes associated with a node

You can select the number of items to display on a page. Valid values are:

- 10
- 20
- 50
- 100

To jump to the next or previous page, or jump to the first or last page, select the arrows at the bottom of the page.

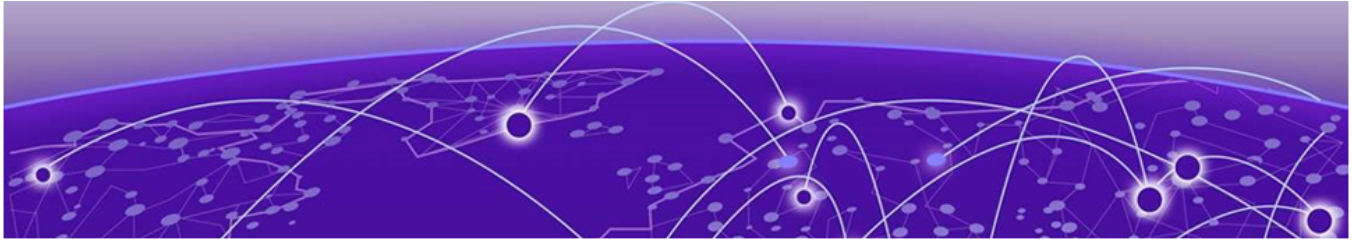

Cluster Settings

The topics in this section describe the options that are available under the **Cluster Settings** menu.

Availability Zones

Before you set up a multi-node cluster, decide on whether to deploy multiple availability zones.

Availability zones let you split a multi-node cluster into separate operational zones where cluster services and applications are distributed across zones. Availability zones add redundancy and improve reliability by ensuring that cluster services and applications remain active even if one of the zones becomes unavailable for any reason.

To deploy multiple availability zones, each zone requires a power supply, cooling, and internet connectivity that is independent of the other zones. You can house the different zones within a single location that has been segregated according to these requirements, or you could add geographic redundancy by housing each zone in a different geographic location.

Feature support includes:

- Maximum of three zones per cluster.
- Minimum of three nodes per zone.
- Each zone within a cluster must have the same number of nodes.
- Individual cluster nodes can belong to a single zone only.

- Minimum cluster size to deploy multiple availability zones is six nodes. This cluster size can provide a two-zone cluster with three nodes in each zone.
- The default cluster setting is a single availability zone with all nodes being located within that zone.

Note
Availability zones can be configured only during the initial cluster creation phase. Once the cluster is created, there is no option to reconfigure the number of zones. You must reinstall and recreate the cluster to change the zone configuration. There is also no option to add or remove availability zones after an upgrade or while adding nodes to an existing cluster.

Example

The following example illustrates a six-node cluster that is split into two availability zones of three nodes each. Each zone has independent power, cooling, and connectivity.

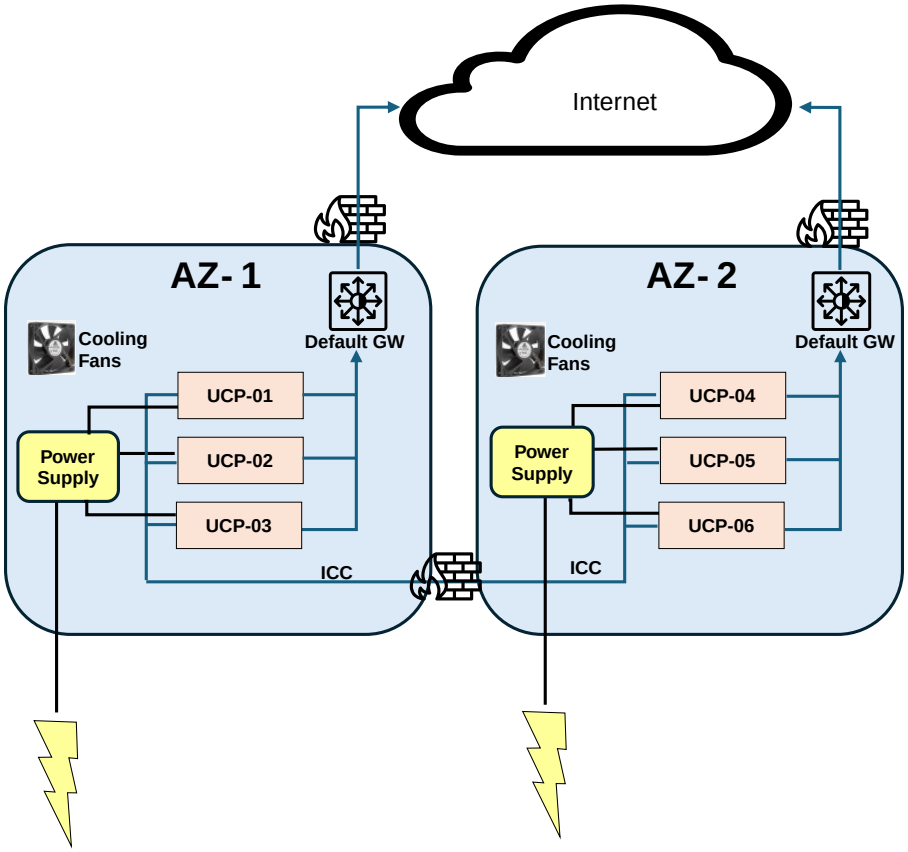

Figure 5: Multiple Availability Zones for a Six-Node Cluster

Cluster Configuration

Go to **Cluster Settings** > **Cluster Configuration** to view the cluster deployment settings and to configure the cluster.

To configure the cluster, complete each of these steps in the cluster configuration process:

1. Select the Deployment Type on page 20
2. Configure Cluster Mode on page 21
3. Configure Pod Network Information on page 21
4. Finish

Note
After the cluster has been set up, the output on the **Cluster Configuration** page changes to provide a read-only view of the cluster configuration.

Select the Deployment Type

From the **Select Deployment Type** field, select the desired deployment type, and then select **Next**.

The following table provides a description of each deployment type option.

Table 4: Deployment Type Options

Deployment Type	Description
ExtremeCloud Edge - Managed Orchestration	ExtremeCloud Edge - Managed Orchestration is a multi-node clustered deployment that offers ExtremeCloud IQ as a distributed cloud application. This deployment provides application delivery and Software as a Service managed by Extreme CloudOps. This deployment supports the following appliances: • 3160C • 4120C-1 For more information, see: • *ExtremeCloud Edge - Managed Orchestration Deployment Guide* • Managed Orchestration Deployment Training Videos
ExtremeCloud Edge - Self-Orchestration	ExtremeCloud Edge - Self-Orchestration deployments offer an integrated Orchestrator for delivery of defined on-premise Universal Container applications. This deployment is offered in standalone deployments of a single node only and supports the following appliances: • 1130C • 2130C • 3150C • 4120C For more information, see *ExtremeCloud Edge - Self-Orchestration Deployment Guide*.

What to do Next

Configure Cluster Mode on page 21.

Configure Cluster Mode

Universal Compute Platform supports a **Standalone** mode and a full **Cluster** mode. Standalone mode requires only one defined node, but a cluster can be deployed with multiple nodes in multiples of three (for example, three, six, nine nodes, or more) depending on your resource requirements.

1. For **Cluster Mode**, select **Standalone** or **Cluster**.
2. (Clusters only). From the **Number of Nodes** drop-down, select the number of nodes for the cluster.
3. (Clusters only). From the **Number of Availability Zones** drop-down, select the number of zones that you want to deploy.
4. Enter the ICC IP Address for each node.
5. Select **Next**.

Note
If you are deploying multiple availability zones, make sure that each ICC IP address falls under the zone where you want to deploy that node.

What to do Next

Configure Pod Network Information on page 21

Configure Pod Network Information

Pods are groups of containers that share networking and storage resources from the same node.

1. Configure the Pod Network configuration settings:
   - Pod Network IP Address (default is 10.96.0.0)
   - Pod Network CIDR (default is 16)
   - Service Network IP Address (default is 10.97.0.0)
   - Service Network CIDR (default is 16)
2. Select **Create Cluster**.

   The cluster is created. If a cluster existed previously, the cluster connections are reset. Then, you must reinstall the engines for each node in the cluster.
3. After the cluster creation process finishes, select **Done**.

Node Additions and Replacements

Add Nodes

Use this procedure to add new nodes to an existing multi-node cluster.

Note

- A multi-node cluster must be a multiple of three nodes. For example, clusters of 3, 6, and 9 nodes are acceptable. As a result, the minimum number of new nodes that you can add is three.
- The Add Nodes feature is not available for standalone cluster deployments.
- If you want to replace a single node that's failed, use the Replace a Node on page 24 procedure.

1. Configure and install the new nodes so that they can be added to the cluster. For each new node, do the following steps. For procedures, see ExtremeCloud Edge - Managed Orchestration Deployment Guide.
   a. Connect to the new node using the Console port.
   b. Run the Basic Configuration Wizard on the new node and assign network settings such as ICC IP addresses, data port IP addresses, DNS, and NTP servers.
   c. Upgrade the software version on the new node so that the new node is running the same version as the rest of the cluster.
   d. Configure VRRP on the new node to match VRRP settings from the existing nodes.
   e. Install the new node on the network so that the node is powered on and has network connectivity to the rest of the cluster. For help, see the Installation Guide for your appliance model.
   f. Repeat these substeps for each new node.
2. Update the cluster configuration to include the new nodes. On the primary cluster node, do the following:
   a. Go to **Cluster Settings** > **Add Node**.
   b. In the **Number of Nodes to Add** drop-down, select the number of nodes that you are adding.
   c. From the **Number of Availability Zones** drop-down, select the number of availability zones to which you are adding nodes.

      Note
      This setting appears only if the existing cluster has multiple availability zones configured.

Note

Keep the following points in mind when adding new nodes to a cluster that is split into multiple availability zones:

- Each availability zone must have the same number of nodes per zone with a minimum of three nodes per zone.
- If the existing cluster has three zones, you must add nodes in multiples of three.
- If the existing cluster has two zones, you must add nodes in multiples of six.
- You cannot reconfigure the number of availability zones in an existing cluster. The cluster must have the same number of zones before and after the node additions.

For more information, see Availability Zones on page 18.

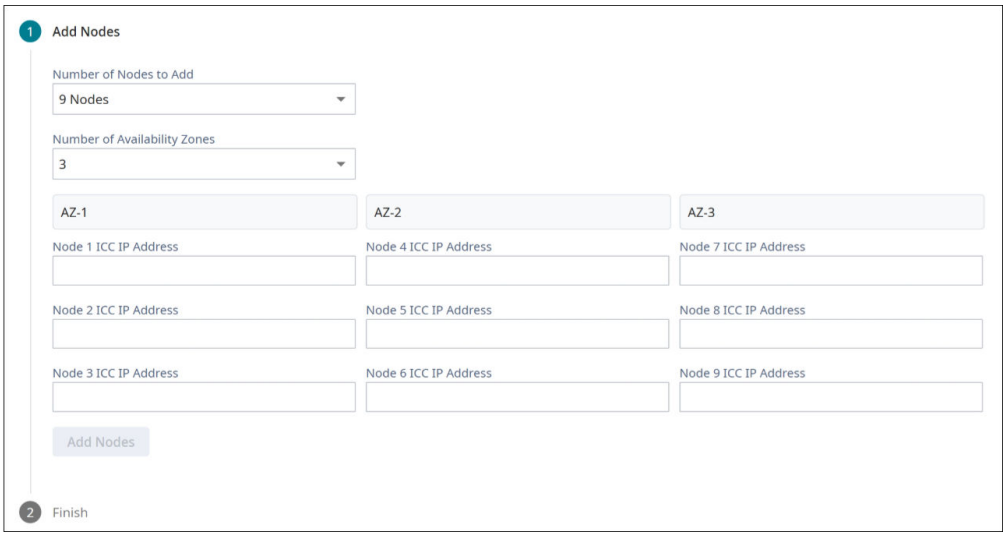

Figure 6: Add Nodes with Multiple Availability Zones

d. Enter the ICC IP Address for each new node and then select **Add Nodes**.
e. Select **OK**.

Replace a Node

Use this procedure if you want to replace a failed node with a new node. The node replacement process removes the failed node from the cluster and inserts the new node in its place.

Note
This procedure applies to multi-node clusters only.

1. For the failed node (the node that's being replaced), do the following:
   a. If you want the new node (the replacement node) to use the same IP settings as the failed node (the node that's being replaced), take a record of the following network settings for the failed node:
      - ICC IP Address—Unlike the other network settings, the new node must use a different ICC IP address than the failed node.
      - Data port IP addresses
      - DNS server addresses
      - NTP server addresses

      Note
      If the failed node is unreachable, you can get the ICC IP address for the failed node from the **Cluster Settings** page of the other cluster nodes.

   b. Remove the failed node from the network so that the node has no connectivity to the other nodes. For example, you could shut the node down or remove the network cables.
2. For the new node (the replacement node), follow the procedures in the ExtremeCloud Edge - Managed Orchestration Deployment Guide to configure the new node. Do the following:
   a. Run the Basic Configuration Wizard on the new node and assign network settings.

      Note
      Make sure to assign an ICC IP address that's different than the ICC IP address of the failed node.

   b. Configure VRRP on the new node to use a lower priority than the other cluster nodes.

      Note
      To ensure that the replacement node successfully joins the cluster, set the VRRP node priority of the replacement node to a value that is lower than the value of the existing nodes. This ensures that the VRRP address is pointing at a working node in the cluster during the joining process. After the replacement node has joined the cluster, you can set the VRRP node priority to first priority if desired, but this is not required.

c. Upgrade the new node to use the same software version that's used by the rest of the cluster.
d. Install the new node on the network so that the node is powered on with connectivity to the rest of the cluster.

3. From the primary node for the existing cluster, start the node replacement process:
   a. Go to **Cluster Settings** > **Node Replacement**.
   b. Select **Next**
   c. For **Select Failed Node**, select the ICC IP address of the failed node.
   d. For **New Node ICC IP Address**, enter the ICC IP address of the replacement node (the new node).
   e. Select **Replace Node**.
   f. Select **OK**.

Engines

The topics in this section describe the options that appear under the **Engines** menu.

Engine Installation Options

From the **Engines** menu, you can install an engine and upgrade an engine application Docker image. The engines that are available to install depend on the deployment type that you selected during the cluster configuration.

To view engine installation information for your deployment, select the deployment type that applies to you:

- ExtremeCloud Edge - Managed Orchestration
- ExtremeCloud Edge - Self-Orchestration

ExtremeCloud Edge - Managed Orchestration

When the deployment type is ExtremeCloud Edge - Managed Orchestration, the only engine that you can install is the ExtremeCloud IQ engine. Complete the following tasks to install the engine.

Table 5: Installation Task Flow for ExtremeCloud Edge - Managed Orchestration

Step	Procedure	Description
1	Run Readiness Assessment on page 27	Optional. Test your system's readiness before you install the engine. Fix any errors in your settings before you install the engine.
2	Install ExtremeCloud IQ Engine on page 28	Install the ExtremeCloud IQ engine.
3	Network Service Configuration on page 28	Configure the mapping of each core service set to the corresponding Virtual Router Redundancy Protocol (VRRP).

Run Readiness Assessment

The Readiness Assessment helps you resolve errors in your network configuration before the ExtremeCloud IQ engine is installed. Run the Readiness assessment prior to onboarding and registering the cluster in Public ExtremeCloud IQ. The cluster registration process automatically notifies CloudOPS and provides basic information on the installation location and network access that is being deployed.

The Readiness Assessment is performed against a specific host at ExtremeNetworks. An assessment service runs that exercises the validation on the access setup through the firewall for the IP Ports that the application(s) require. The assessment services are installed at `ucp0-console.extremecloudiq.com`.

The assessment does the following:

- Pulls service groups and ports for inbound and outbound connections.
- Lets you enter the IP addresses that you plan to deploy.
- Tests your configuration and reports the results using a PASS and FAIL convention.

Note
Make sure that your firewall is configured to allow external and inbound access in relation to the firewall rules and service sets that appear in this document to ensure that the test succeeds.

Note
Make sure that your firewall is configured to allow external and inbound access in relation to the firewall rules and service sets that appear in the *ExtremeCloud Edge - Managed Orchestration Deployment Guide* to ensure that the test succeeds.

1. Go to **Engines** > **Installation**.
2. From the ExtremeCloud IQ pane, select **Readiness Assessment**.
3. When prompted, enter the **VRRP IP Address** and **External IP Address** that you plan to deploy for each service group and port. See the subsequent table for more information on these fields.
4. Select **Test**.
5. For any tests that received a FAIL result, or for any other error message, make the required configuration corrections and rerun the test.
6. If you receive a PASS for all checks, proceed to engine installation.

The following table provides information on the fields that display around the Readiness Assessment.

Table 6: Readiness Assessment Field Descriptions

Field	Description
Outbound	
Port	The port over which the outbound connection is tested.
Protocol	The protocol that is in use for outbound connections on this port.

Table 6: Readiness Assessment Field Descriptions (continued)

Field	Description
Result	The result of the test. Possible results include: • PASS • FAIL
Error	For tests that fail, the value in this field provides information about the problem so that you can fix it.
Inbound	
Service Group Name	The name of the service group (or service set) that accepts incoming connections to this external IP address.
Port	The port over which the inbound connection is tested.
Port Name	The name of the port.
Protocol	The protocol that is in use for inbound connections to this port and external IP address.
VRRP IP Address	The internal VRRP IP address that provides load balancing and high availability for inbound connections to this service group.
External IP Address	The public IP address that accepts incoming connections for this service group. The connection is port-forwarded to the internal VRRP IP address for this service group.

Install ExtremeCloud IQ Engine

Install ExtremeCloud IQ engine once from a single node.

To install an engine instance:

1. Go to **Engines**.
2. From the ExtremeCloud IQ pane, select **Install**.

   After installation is complete, a confirmation notice is displayed and the XIQ instance displays.

Network Service Configuration

The **Network Service Configuration** tab displays the mapping of each core service set to the corresponding Virtual Router Redundancy Protocol (VRRP).

Note
Network registration is configured during the initial Universal Compute Platform setup process. For complete instructions on registering a network account, see the *ExtremeCloud Edge - Managed Orchestration Deployment Guide*.

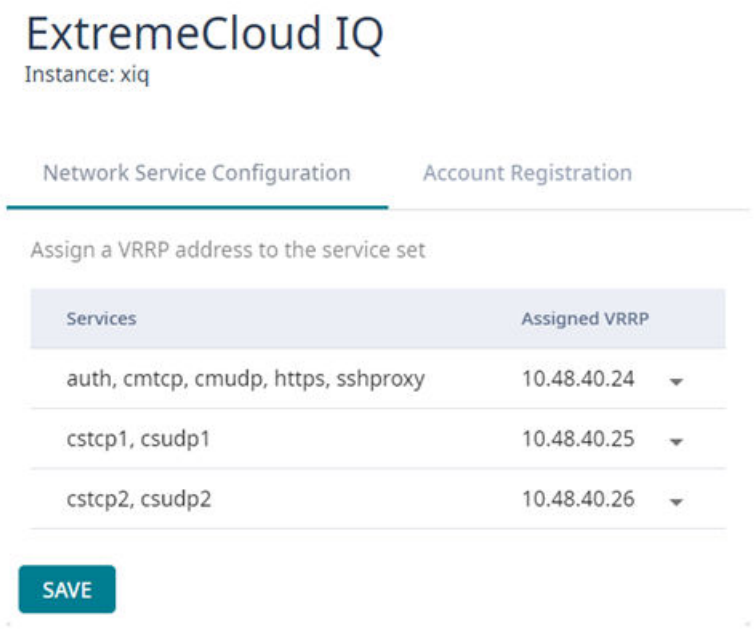

Figure 7: ExtremeCloud IQ Network Service Configuration Details

Account Registration

Note
The **Account Registration** tab is used with legacy deployments of Distributed Cloud only. For onboarding information that is related to ExtremeCloud Edge – Managed Orchestration, see the *ExtremeCloud Edge - Managed Orchestration Deployment Guide for Universal Compute Platform*.

Create an ExtremeCloud IQ user account through Universal Compute Platform. Go to **Engines** > **Account Registration** and fill out the form in Figure 8. Then, select **Register**.

You will receive an email confirming your registration.

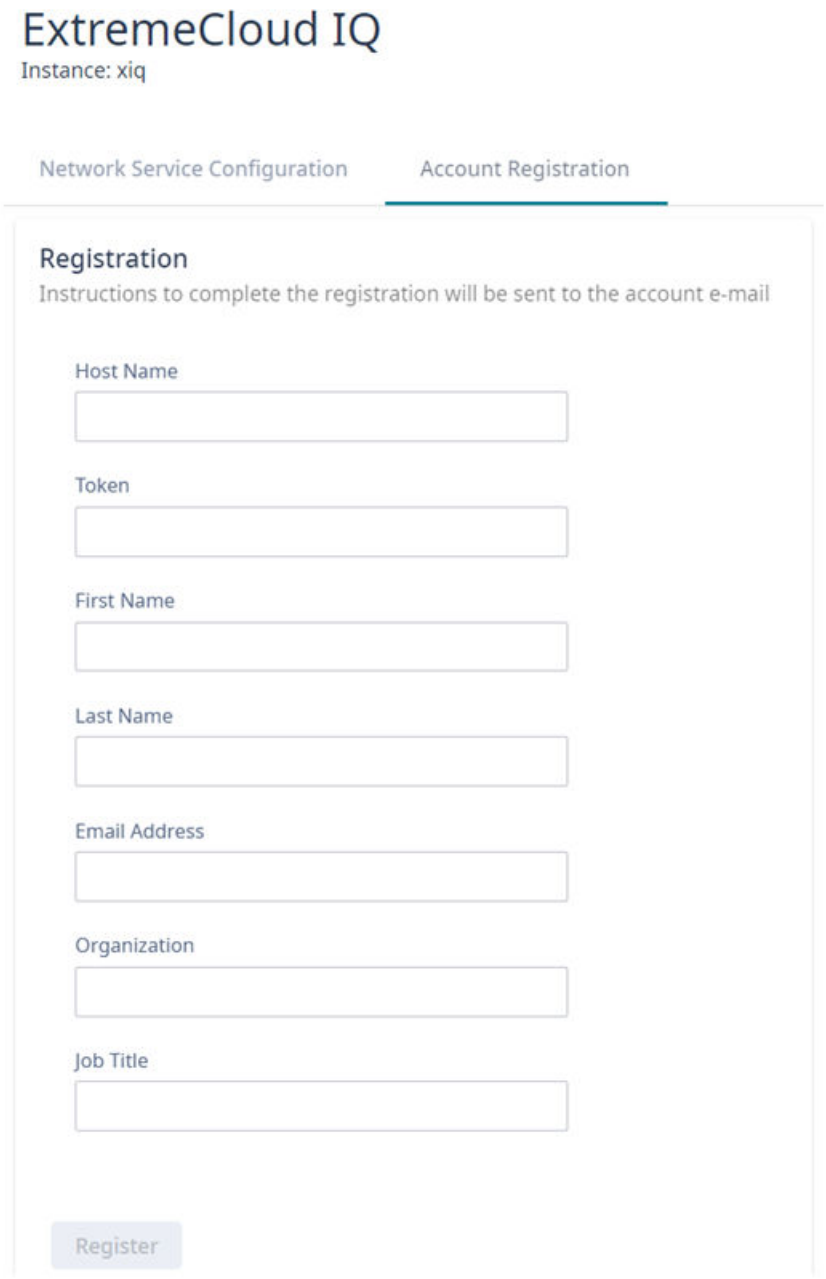

Figure 8: ExtremeCloud IQ Account Registration Form

ExtremeCloud Edge - Self-Orchestration

When the Deployment Type is ExtremeCloud Edge - Self-Orchestration, you have a variety of engine options to choose from, including the following:

- ExtremeWireless WiNG™ Controller
- Extreme Tunnel Concentrator
- ExtremeCloud™ IQ Controller

Install an Engine

To install an engine for an ExtremeCloud Edge - Self-Orchestration deployment, complete the following tasks in order.

Table 7: Initial Engine Application Installation

Step	Procedure	Description
1	Download Docker Application Image on page 31	Download the Extreme docker application Image file from the support portal.
2	Upload Docker Application Image on page 31	Upload the docker application image to Universal Compute Platform.
3	Install Engine Application on page 32	Install the application engine on Universal Compute Platform.
4	Deploy Application Image on page 32	Deploy the Extreme application image file
5	Configure Interface Settings for Engine Application on page 33	Copy the application's Locking ID (for licensing) and, if required, assign a VRRP IP alias.

What to do Next

- If you need to upgrade application firmware, go to Engine Upgrades on page 35.
- If you want to onboard your deployment to the cloud, following cloud onboarding procedures in ExtremeCloud Edge - Self-Orchestration Deployment Guide.

Download Docker Application Image

Download the application Docker image file from the Extreme Networks support portal.

To obtain the Docker image file, go to the Extreme Networks support portal to download the application Docker image.

For example, from the ExtremeWireless WiNG™ product page, download `cx-9000.tar`

Upload Docker Application Image

Upload the engine application Docker image to Universal Compute Platform.

1. Go to **Engines** > **Image Management**.

2. Complete one of the following options:
   - Select the **Choose Image File** pane, then navigate to the image file and select it.
   - Drag and drop the image file onto the **Image File** pane.

   The uploaded image file displays below the **Choose Image File** pane along with other uploaded image files.

> **Note**
> You can also delete an uploaded image file. From the **Image Management** page, select the check box next to the image file, and select (Delete). To refresh the image file list, select (Refresh).

Install Engine Application

To install the engine application, take the following steps:

1. Go to **Engines** > **Installation**
2. From the pane for the application that you want to install, select **Install**.

> **Note**
> - If you have not yet uploaded the application docker image file, you will be prompted to do so.
> - The installation time depends on a variety of factors. Be prepared for it to take some time.

A confirmation notice displays after the installation completes. Only one instance is required for the cluster.

Deploy Application Image

After you have uploaded the application image file and installed the application Docker image, deploy the application to a node.

1. Go to **Engines** > **Installation**.
2. Select the engine instance link. For example, "cx9000 #1".
3. Select **Deploy**.
4. Save your changes.

Configure Interface Settings for Engine Application

Use this procedure to configure the final application interface settings immediately ater you deploy an application. You can use these settings to assign a VRRP IP alias to the application (which is required for some applications).

Note

- Refer to the documentation for your engine application for application-specific requirements. Configuration requirements can vary per application.
- For more information on the fields that appear on this page, see Engine Application Settings on page 33.

1. Go to the application interface settings page for your application.

   The page launches automatically after you select the **Deploy** option for a newly installed application. Otherwise, go to **Engines** > **Installation** and then select the link for the engine instance.
2. Take a copy of the **Locking ID** (Serial Number). Most applications require you to provide this value during license activation.
3. Optional. If your engine application requires you to configure a VRRP IP alias:
   a. Select the **Network Service Configuration** tab.
   b. From the **Assigned Virtual IP Address** field, select a VRRP address that you assigned previously on a data port.
   c. Select **Save**.
4. Optional. To launch the application user interface, select the **Instance Web Interface** link.

Engine Application Settings

After you deploy a newly installed engine, the settings page for the installed application displays with settings for that engine instance. You can also access this page by going to **Engines** > **Installation** and then selecting the instance. The name of the application displays at the top of the screeen (for example, "**Extreme Tunnel Concentrator**").

Depending on which application engine you installed, the displayed fields may vary:

- Image—The image name.
- Version—Version number.
- Instance—Name of the node instance.
- Status—Status of the application engine. For example, **Running**.
- Hostname—Hostname of the application engine instance.
- OS Version—Operating system of the application engine.
- Locking ID—Locking ID (or Serial Number) of the application engine instance. Copy this number as most applications require you to enter it during license activation.
- Node—The name of the host node.
- Port Information—Port information for the host node
- Instance web interface—The IP address to access the installed application. Select this link to launch the application's UI. If the application requires a VRRP IP alias, you

must first go to the **Network Service Configuration** tab and set an **Assigned Virtual IP Address** first.

The following four tabs display across the middle of the page. You can use these tabs for diagnostic information about application services.

Network Service Configuration

For applications that require a VRRP IP alias, this tab displays the **Assigned Virtual IP Address** field. If the application requires that you assign a VRRP IP alias, from the drop-down , select a VRRP address that was configured previously on one of the data ports. The VRRP alias provides an IP address for access to the application's management interface.

VRRP enables a virtual router to act as the default network gateway, improving host network reliability and performance.

Note
For a given VRRP address to display, the VRRP address must have been added already on one of the data ports for the Universal Compute Plaform host.

Statistics

Compute statistics and node drive volume statistics are available for CPU usage and memory usage.

Logs

A log file is available for each node instance. Log entries include the following:

- Timestamp of log entry
- System Component
- Message log level
- Message content

Console

A live console is available from each engine instance for diagnostics and troubleshooting. To open a live console and connect to a container or virtual machine instance (VMI), from the engine **Console** tab, select **Attach**.

Image Management

From the **Engines** > **Image Management** page, you can manage uploaded engine images for Universal Compute Platform.

- To upload a new image to Universal Compute Platform, complete either of the following steps:
  - Select the **Choose Image File** pane, browse to the image file and select it.
  - Drag the image file from a local drive and drop it onto the Universal Compute Platform desktop.
- To delete an existing image file from Universal Compute Platform, select the existing image and select (Delete).

- To refresh the list of uploaded images, select  (Refresh).

Engine Upgrades

Universal Compute Platform has multiple methods for upgrading container applications. Select the upgrade method that fits your application type:

- **Self-Orchestrated applications**—For self-orchestrated applications that support external upgrades, see Upgrade an Application (Self-Orchestrated) on page 35.
- **Applications with built-in upgrade functionality**—For applications with built-in upgrade functionality, you can upgrade from the application interface. Refer to the application documentation for details.
- **Applications that do not support either upgrade method**—For these applications, uninstall the current image and then install the new image. Note that this method requires you to reconfigure your settings.

Upgrade an Application (Self-Orchestrated)

Use this procedure to upgrade a self-orchestrated engine application from the Universal Compute Platform user interface. This procedure upgrades the application while retaining existing settings.

Note
You must have the new application image file. For Extreme Networks applications, download the install image from the *Extreme Networks Support Portal* and save it to a local drive.

1. Log in to the Universal Compute Platform interface.
2. Upload the new application image file:
   a. Go to **Engines** > **Image Management**.
      A list of uploaded images displays under the **Choose Image File** pane.
   b. To upload the new image, complete either of the following steps:
      - Select **Choose Image File**, then browse to the image file and select it. Or,
      - Drag the image from your local drive and drop it on the **Choose Image File** pane.

      Note
      To delete an image file, select the check box next to the image and select  .

3. Upgrade the application:
   a. Go to **Engines** > **Installation**.
   b. Select the application instance that you want to upgrade.

c. Select **Upgrade application**.
d. Select **OK**.

   Universal Compute Platform creates a new container with the upgraded application image and existing settings. The old container is terminated.

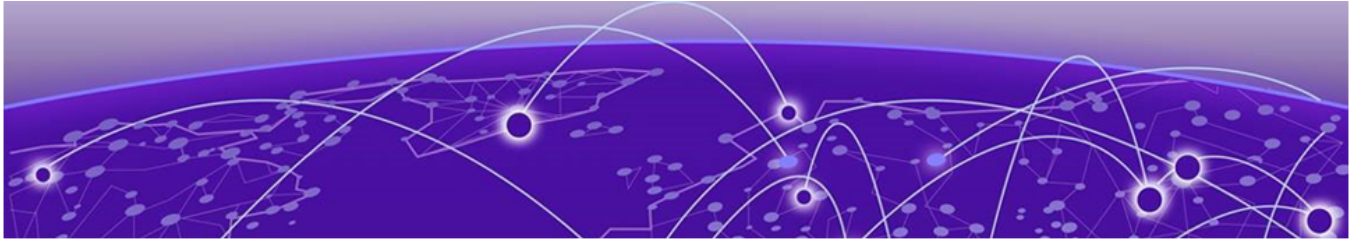

Tools

The topics in this section describe the settings that appear under the **Tools** menu.

Logs

Universal Compute Platform offers logs to help you understand and troubleshoot the network. Go to **Tools** > **Logs** and select the tab with the log type that you want to view. For log details, refer to Table 8 on page 37. The log options are:

- **Events**
- **Audit Log**
- **Access Events**
- **Hardware Events**—Not available on all Universal Compute Platform models.

To filter the list of logs by date range, select the appropriate Start Date and End Date from the filter option. To clear the filter, select **Reset**.

Table 8: Log Types

Log Type	Details
Events	To view a list of network events, select the **Events** tab. The following information displays for each event: • **Time**—Time the event occurred. • **Type**—The type of event: Info, Minor, Major, or Critical. • **Component**—The component of Universal Compute Platform that was affected. For example, Rest API or Startup Manager. • **Description**—A description of the event.
Audit Log	To view the Audit Log, select the **Audit Logs** tab. The following information displays in audit logs: • **Time**—The time that the logged item occurred. • **Username**—The username of the system administrator. • **Context**—The context for the event. • **Description**—A description of the logged item.

Table 8: Log Types (continued)

Log Type	Details
Access Events	To view access events logs, select the **Access Events** tab. This log provides information on access events, including authentication attempts, the opening and closing of sessions, and a record of blocked access if the number of failed logins exceeds the allowed threshold. The page displays information under the following categories: • **Time**—The time the event occurred. • **Source**—The source of the event. • **Message**—The event message.
Hardware Events	To view hardware events logs, select the **Hardware Events** tab. These logs provide current status and alert information on hardware items such as the power supply, fans, CPUs, and memory storage. The page displays information under the following categories: • **Timestamp**—The time that the event occurred. • **Source**—The source of the event. For example, Power Unit Pwr Unit Status. • **Message**—The event message. For example, Power off/down. • **Status**—The current status. For example, Asserted or Deasserted.

You can select the number of items to display on a page. Valid values are:

- 10
- 20
- 50
- 100

To jump to the next or previous page, or jump to the first or last page, select the arrows at the bottom of the page.

Diagnostics

Go to **Tools** > **Diagnostics** for tools that help you troubleshoot your network. You can use the following tools:

- **Utilities**—Use **Ping** or **Trace Route** to troubleshoot specific IPs and FQDNs
- **TCP Dump Management**—Run file captures on specific interfaces.

Utilities

Use wireless controller utilities to test a connection to the target IP address (or Fully-Qualified Domain Name) and record the route through the Internet between your computer and the target address. You can also use controller utilities to capture

exception traffic, which can be useful for network administrators when debugging network problems.

Table 9: Network Utilities

Field	Description
Target IP Address or Fully-Qualified Domain Name (FQDN)	IP address or FQDN for the test target.
Use specific source interface	Indicates if a specific interface will be selected for the test. Select the interface from the Select Interface field. When this option is cleared, Universal Compute Platform runs the test based on the interface selected in the routing table.
Select Interface	Used with Specific Source Interface option. See list of possible interfaces on the **Interface** tab.
Ping	Initiate the Ping network utility to determine reachability of the IP address or FQDN that you specify.
Trace Route	Initiate the Trace route command, which traces the path of a packet from Universal Compute Platform to the IP address or FQDN that you specify. It lists the routers it passes until it reaches its destination, or fails to. It also indicates the length of each hop.

TCP Dump Management

The following table describes the fields in the TCP Dump Management section:

Table 10: TCP Dump Management

Field	Description
Interface	Target interface. See the list of possible interfaces on the **Interface** tab.
Filename	Specify the name of the dump file.
Save File To	Specify where to save the dump file.
Capture File Size (MB)	Specify the maximum limit of the dump file in MB. This feature enables you to control the size of the resulting dump file so the file does not become too large.
Capture Files	List of previously created dump files. Select a file to take action.

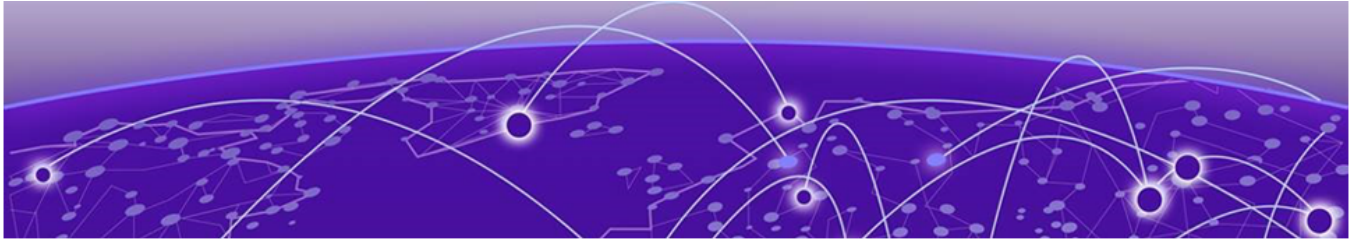

Administration

The topics in this section describe the settings that appear under the **Administration** menu.

Manage User Accounts

This topic outlines how to manage user accounts on the Universal Compute Platform controller. For information about registering for an ExtremeCloud IQ user account, see Account Registration on page 29.

Universal Compute Platform offers the following levels of user access on the controller:

- Full Admin
- Read Only

Full Administrators can create and manage controller user accounts. This guide outlines the following procedures:

- Add new accounts
- Modify account settings
- Delete user accounts

For information on the settings that you can configure for a user account, see Account Settings on page 42.

Add a User Account

To add a user account:

1. Go to **Administration** > **Accounts**.
2. Select **New Account**.

3. Configure the account settings.

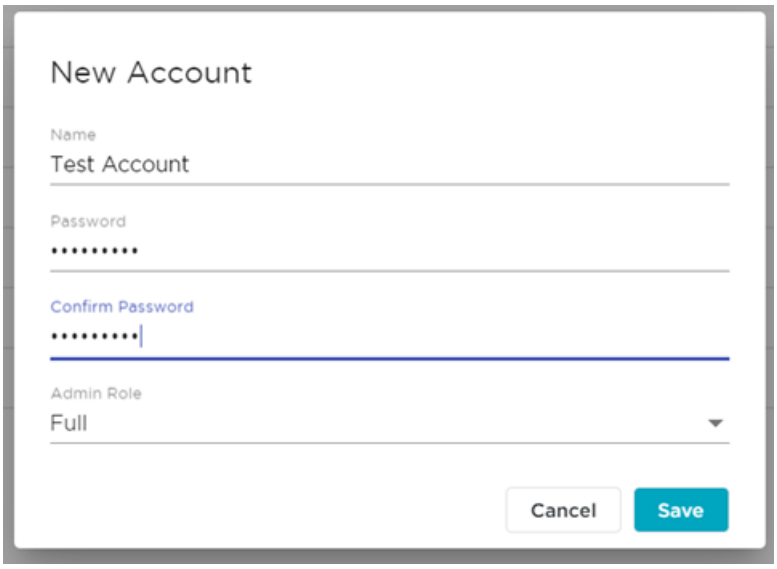

Figure 9: Create New Account

Modify a User Account

To modify a user account:

1. Go to **Administration** > **Accounts**.
2. Select ✎ next to the account that you want to modify.
3. Select **Change Password**.

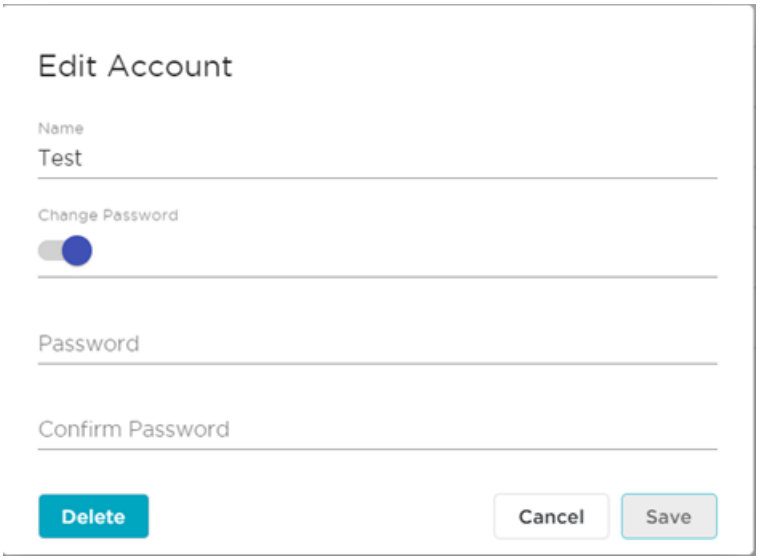

Figure 10: Edit Account Details Dialog

4. In the Password field, enter a password.
5. In the Confirm Password field, enter the same password again.

6. Select **Save**.

Note
For more information on user account settings, see Account Settings on page 42.

Delete a User Account

To delete a user account:

1. Go to **Administration** > **Accounts**.
2. Select  next to the account that you want to delete.

   The **Account Settings** dialog opens.
3. Select **Delete**.

   A confirmation dialog displays.
4. Select **OK** to confirm that you want to delete the account.

Account Settings

Configure the following user account settings:

Name

Name for the user account.

Password

Password for the user account. The password must be between 8 and 24 characters.

Confirm Password

Enter the password for the user account a second time.

Admin Role

The access level for the user account. Valid values are:

- Full Admin
- Read Only

System Configuration

System administrators can do the following from the **Administration** > **System** menu:

- Configure network interfaces and network time
- Manage Universal Compute Platform upgrades and system maintenance
- Configure availability mode for network failover and redundancy
- View system logs and information.

Configuration

Go to **Administration** > **System** > **Configuration** to complete a configuration backup or restore.

This backup and restore procedure is limited to configuration files and, optionally, logs and audit files. A system backup is a different procedure. A system backup is a full system snapshot rescue file (*-rescue-user.tgz). Creating a full system rescue file is an option during the system upgrade process.

Before you perform a backup procedure, decide what to back up and where to save the backup file:

Before you perform a backup procedure, decide what to back up and where to save the backup file:

- Select back up configs, logs, and audit or back up configuration only.
- Select a location to store the backup file.
- Select **Local** as the backup location.
- (Optional) Configure a backup schedule.

Note
It is a best practice to set up a scheduled backup for all managed appliances.

On-demand backups can only be stored locally, while scheduled backups can be stored on a mounted flash drive or on a remote server.

You can select from the following tasks:

- Run a Backup on page 43
- Restore Backup File on page 43
- Schedule a Backup on page 45
- Configure SNMP on page 46

Run a Backup

Use this procedure to run an on-demand configuration backup. A configuration backup is limited to configuration files and, optionally, logs and audit files.

1. Go to **Administration** > **System** > **Configuration**.
2. Select the **Backup/Restore** tab.
3. Select **What to back up**. The options are:
   - **Configs, Logs and Audit**
   - **Configurations only**
4. Select **Where to Backup** (for example, **Local**).
5. Select **Start Backup**.

Restore Backup File

Use this procedure to restore the appliance from a selected configuration backup file.

1. Go to **Administration** > **System** > **Configuration**.
2. Select the **Backup/Restore** tab.

3. If no backup file displays under **Select Backup**, then complete the following steps to upload the backup file:
   a. Under **Restore**, select ⚙, then select **Upload**.
   b. Select the **Upload Method**. Valid values are:
      - **HTTP**
      - **FTP**
      - **SCP**
   c. If you chose HTTP as the upload method, do either of the following and then select **Close**:
      - Select **Choose Backup file**, then browse to the backup file and select it
      - Use the cursor to drag and drop the backup file from a local drive onto the desktop.

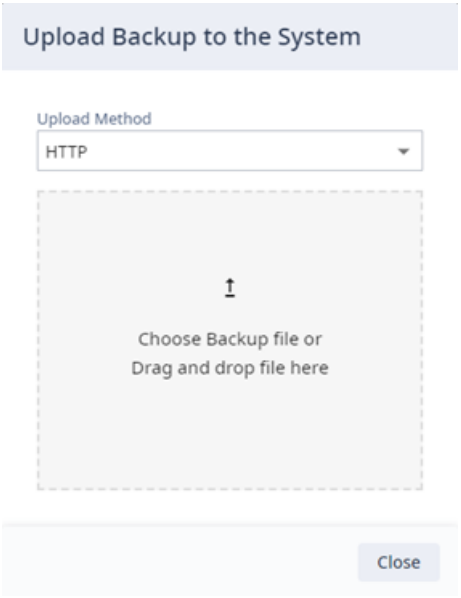

Figure 11: Upload Controller Image

   d. If you chose FTP or SCP, enter the remore server details and then select **OK**:
      - **Server IP**—Enter the IPv4 address of the remote server.
      - **Username**—Enter a username that has access to the remote server.
      - **Password**—Enter the password that authenticates the username on the remote server.
      - **Directory**—Enter the directory where the backup file is saved.
      - **Filename**—Enter the filename of the backup file.
      - **Destination**—Enter the upload destination. For example, Local.
4. From the **Select Backup** drop-down, select the uploaded backup file.
5. Run the restore job.

Schedule a Backup

When you schedule a backup, you can choose to upload the backup to a server or have the scheduled backup saved locally or on an external flash drive. To schedule a backup:

1. Go to **Administration** > **System** > **Configuration** > **Schedule Backups**.

   The **Schedule Backup** dialog displays.

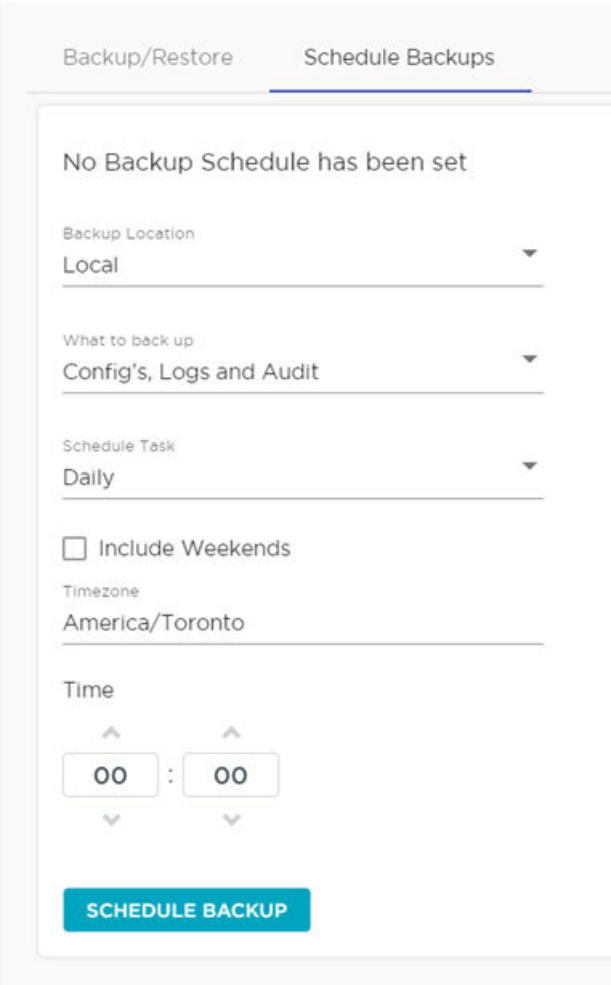

Figure 12: Schedule Backup Dialog

2. Configure the following parameters:

   Backup Location

   Indicates where to send the backup file. Valid values are: Local or Remote. When sending a backup to a remote server, configure the server properties.

   What to back up

   Indicates the content of the backup file. Valid values are: Configs, Logs and Audit (which is a full backup), or Configuration files only.

   Schedule Task

   Indicates when the backup task runs. Valid values are: Daily, Weekly, Monthly.

Include Weekends

Select this check box to include weekends in the backup schedule.

Timezone

Select the timezone to be used for the backup.

Time

Set the time of day for the scheduled backup.

3. If you selected **Remote** for the **Backup Location**, enter the remote server details:

Protocol

Select FTP or SCP.

Server IP

Enter the IPv4 address of the remote server.

Username

Enter a username with login access to the remote server.

Password

Enter a password that authenticates the username on the remote server.

Directory

Enter the directory where the backup file is to be saved.

4. Select **Schedule Backup**.

Configure SNMP

Use this procedure to configure Universal Compute Platform as an SNMP agent. You can use SNMPv2c or SNMPv3 to monitor the system using a Network Management Station (NMS) server. Unless stated otherwise, all steps apply to both SNMP versions.

Note
For detailed information on the SNMP fields and configuration options, see SNMP Settings on page 48.

1. Go to **Administration** > **System** > **Configuration**.
2. Select the **SNMP** tab.
3. Select the **SNMP Mode** from the following options:
   - **Disabled**—To disable SNMP, select this mode and then select **Submit**.
   - **SNMPv2c**—Select this mode to enable SNMPv2c. This mode includes default public and private communities. You can use the defaults or assign new settings.
   - **SNMPv3**—Select this mode to enable SNMPv3. This mode includes an autogenerated **Engine ID**. You can use the default ID or assign a new value.
4. (SNMPv3 only). Enter a **Context String**.
5. (SNMPv3 only). Configure **SNMP Users**:
   a. Select **Add User**.
   b. Enter a **Name** for the user.

c. Select a **Security** level for SNMP messaging for this user:
   - Authentication, Privacy
   - Authentication, No Privacy
   - No Authentication, No Privacy

d. If the selected security level includes authentication, select an **Auth Type** and enter an **Auth Password**.
e. If the selected security level includes privacy, select a **Privacy Type** and enter a **Privacy Password**.
f. Select **Save** to save the user configuration.

6. Configure up to two SNMP notifications:
   a. Select **Add Notification**.
   b. Select a **Name** for the notification.
   c. For **Server Address**, enter the IP address of the NMS server.
   d. For **Server Port**, enter the receiving port on the NMS server.
   e. Select **Save** to save the notification.
7. From the **Forward Traps** tab, select the severity threshold for the sending of SNMP messages to the NMS server.
8. Select **Submit** to save the SNMP configuration.

Note
To edit settings for an existing SNMP community, user, or notification, or to update an authentication or privacy password for an existing user, select the adjacent icon, configure the new settings, and select **Save**.

Note
Optionally, you can download a MIB resource package zip file from the Extreme Networks Support Portal. You can upload the MIBs to an NMS server and track application status using SNMP.

SNMP Settings

To access SNMP Settings, go to **Administration** > **System** > **Configuration** and select the **SNMP** tab. For field descriptions of the available settings, see Table 11. Note that each setting applies to both SNMP versions, unless noted otherwise.

Table 11: SNMP Settings

Field	Description
SNMP Mode	Select the SNMP mode from the following options: • Disabled—SNMP is disabled. • SNMPv2c—SNMP is enabled using SNMPv2c. • SNMPv3—SNMP is enabled using SNMPv3.
SNMP Communities	(SNMPv2c only) This section contains SNMP community settings for SNMPv2c. There is a limit of two SNMP communities. SNMPv2c provides the following two default SNMP communities. You can use these default communities, edit them, or delete them and then add your own. • public—includes READ-LEVEL access. • private—includes WRITE-LEVEL access. From the **Actions** column, select the following icons to: • —Select this icon to edit settings for the adjacent community. • —Select this icon to delete the adjacent community. To add a new community, select **Add Community**, configure the following fields, and then select **Save**. • **Name**—Enter a name (between 1 and 32 characters). • **Access Level**— Select READ-LEVEL or WRITE_LEVEL.
Context String	(SNMPv3 only) Enter a text string between 1 and 32 characters to identify the SNMP configuration.
Engine ID	(SNMPv3 only) The engine ID for the SNMP agent. SNMPv3 auto-generates an engine ID, but you can assign a new value. The ID must be between 1 and 27 characters, and includes support for special characters, but does not include support for the following characters: < | &.

Table 11: SNMP Settings (continued)

Field	Description
SNMP Users	(SNMPv3 only) This section contains settings for SNMPv3 users. To add a new SNMP user, select **Add User**, configure the following user settings, and then select **Save**: • **Name**—A text string between 1 and 32 characters to identify the user. • **Security**—The security level for SNMP messaging. The options are: ◦ Authentication, Privacy (AUTH_PRIV)—Messaging is authenticated and encrypted. ◦ Authentication, No Privacy (AUTH_NO_PRIV)—Messaging is authenticated but non-encrypted. ◦ No Authentication, No Privacy (NOAUTH_NOPRIV)—Messaging is non-authenticated and non-encrypted. • **Auth Type**—If the security level includes authentication, this field specifies the authentication hash algorithm. The options are: MD5, SHA, SHA224, SHA256, SHA384, and SHA512. In addition, when you add or edit users with authentication, assign the following authentication fields: ◦ **Auth Password**—An authentication password between 8 and 32 characters with alpha, numeric, and special characters. ◦ **Mask**—When checked (the default setting), the Auth Password cannot be read in the GUI. Otherwise, the password is visible. • **Privacy Type**—If the security level includes privacy (encryption), this field specifies the encryption algorithm. The options are: DES, AES, AES192, and AES256. In addition, when you add or edit users with privacy, assign the following prvacy fields: ◦ **Privacy Password**—An encryption password between 8 and 32 characters with alpha, numeric, and special characters. ◦ **Mask**—When checked (the default setting), the Privacy Password cannot be read in the GUI. Otherwise, the password is visible. For existing SNMP users, use the **Actions** menu to update settings: • —Select this icon to edit settings or authentication and privacy passwords for the adjacent user. Assign new settings and select **Save**. • —Select this icon to delete the adjacent SNMP user.

Table 11: SNMP Settings (continued)

Field	Description
SNMP Notifications	This set of fields lists SNMP notifications that SNMP watches for. You can add up to two SNMP notifications. To add a new notification, select **Add Notification**, set the following fields, and select **Save**. • **Name**—Select the name for the notification. With SNMPv2c, the version displays in this field. With SNMPv3, select an SNMP user. • **Version** (SNMPv3 only)—The SNMP version for this notification. • **Server Address**—The IP address of the Network Management Station (NMS) server to which you are sending this SNMP notification. • **Server Port**—The receiving port on the NMS server. For **Actions**, you can select from the following: • —Select this icon to edit the adjacent notification. • —Select this icon to delete the adjacent notification.
Forward Traps	From the drop-down, select the severity threshold for sending event logs to the SNMP server as forward traps. Logs with a severity that meets or exceeds this threshold get forwarded to the NMS server as traps. The options, in order of least-severe-to-most severe, are as follows: • Informational • Minor • Major • Critical (default setting)

System Logging

Syslog event reporting uses the syslog protocol to relay event messages to a centralized event server on the enterprise network. In the protocol, a device generates messages, a relay receives and forwards the messages, and a syslog server receives the messages.

System Log Level

Determines the error severity that is logged for the appliance. Select the least severe log level that you want to receive: Information, Minor, Major, Critical. For example, if you select Minor, you receive all Minor, Major and Critical messages. If you select Major you receive all Major and Critical messages. The default is Minor.

Syslog

Provide the IP Address of 1-3 syslog servers and enable the type of messages that you want to send to the syslog servers.

- **Send all Service Messages**
- **Send Audit Messages**

Note
To synchronize the logs, the syslog daemon must be running on both the appliance and on the remote syslog server. When you change the log level on the appliance, you must modify the appropriate setting in the syslog configuration on remote syslog server.

Facility Codes

Facility codes identify log streams in the remote syslog server. Select a unique facility code (local.0 - local.6) for each Universal Compute Platform facility to differentiate the log streams and facilitate the filtering of messages.

The facility code applies to all servers. Select a facility code for each of the following:

- Application Facility
- Service Facility
- Audit Facility

Maintenance

Perform cluster maintenance and tech support from the **Maintenance** page. Go to **Administration** > **System** > **Maintenance** .

The following table describes the commands and options that are available.

Table 12: Maintenance Fields

Field	Description
System Actions	
Reset Configuration	This command does the following: • Resets all user configurations • Provides the option to reset the ICC (management) port configuration • Resets the Kubernetes node • Resets the shared file system **Note:** Do not reset the configuration for single nodes within a multi-node cluster unless you plan to reset the configuration for all nodes in the cluster. Running this command on a single node only causes an invalid cluster configuration. If you need to replace a single node that's failed, use the procedure Replace a Node on page 24.
Reboot	The Universal Compute Appliance shuts down, then reboots. A warning message is displayed, asking you to confirm your selection.
Shut Down	The system enters the halted state, which stops all functional services and the application. A warning message is displayed, asking you to confirm your selection. To restart the system, the power to the system must be reset.
Cluster Actions	
Reset Node	This command does the following: • Resets the Kubernetes node • Resets the shared file system **Note:** Do not reset single nodes within a multi-node cluster unless you plan to reset all nodes in the cluster. Resetting a single node only causes an invalid cluster configuration. If you need to replace a single node that's failed, use the procedure Replace a Node on page 24.

Table 12: Maintenance Fields (continued)

Field	Description
Session The Web Session Timeout determines the web session inactive window before the session times out. Enter the value as hours : minutes. The range is 1 minute to 168 hours (7 days).	
Web Session Timeout (Hours)	The number of hours that you want to add to the web session inactive timer. Valid values are 0-168 hours.
Web Session Timeout (Minutes)	The number of minutes that you want to add to the web session inactive timer. Valid values are 0-59 minutes.
Tech Support	
Generate Tech Support	When generating a tech support file for troubleshooting, first select one of the following file criteria: • Appliance • Log • All (default value) To generate a tech support file, select **Generate Tech Support**. To download the generated file, select  .

Network Setup

To view the network setup, go to **Administration** > **System** > **Network Setup**. where you can view and edit network settings within the following categories:

- Host Attributes
- L2 Ports
- ICC Interfaces
- Interfaces
- Static Routes

Host Attributes

The **Host Attributes** section displays the fields in the following table.

Table 13: Host Attribute Field Descriptions

Host Attribute Setting	Description
Host Name	The host name of the appliance. Do not edit this setting after initial deployment.
Domain Name	The domain name of the appliance. Do not edit this setting after initial deployment.
Default Gateway	The Default Gateway IP address is the global default IP route setting for the appliance. Valid values are: • The admin topology gateway address • Any IP address on the physical Interfaces or Bridge at AC VLAN topology subnets.

Table 13: Host Attribute Field Descriptions (continued)

Host Attribute Setting	Description
DNS Server 1	The IP address of the primary DNS server for the appliance.
DNS Server 2	The IP address of the secondary DNS server for the appliance. A secondary DNS server is not mandatory, but provides DNS redundancy.

You can edit the Default Gateway, DNS Server 1, and DNS Server 2 settings right from the Network Setup page by doing the following:

1. Go to **Administration** > **System** > **Network Setup**.
2. Under **Host Attributes**, enter the new settings for the applicable fields.
3. Select **Save**.

L2 Ports

The **L2 Ports** section displays the fields the following table. Use the L2 Ports information to understand the OSI Layer 2 (Data Link Layer) physical topology of the data plane. These ports represent the actual Ethernet ports.

Table 14: L2 Port Field Descriptions

L2 Port Settings	Description
Status	The link status. The status may show green (active) or red (inactive).
Name	Name of the L2 port.
Speed	Connection speed of the L2 port
VLANs	Assigned VLANs on the port.

Table 14: L2 Port Field Descriptions (continued)

L2 Port Settings	Description
Member Ports of LAG	Shows whether LAG is enabled or disabled on the port. You can toggle LAG on or off.
Stats	Select the link to open a popup that displays more information on the port. Depending on the type of port, the popup may display: • MAC address of the port • Port speed (in bits per second) • Frames sent and received (depending on port and traffic, this figure may be split by unicast, multicast, and broadcast traffic) • Octets sent and received 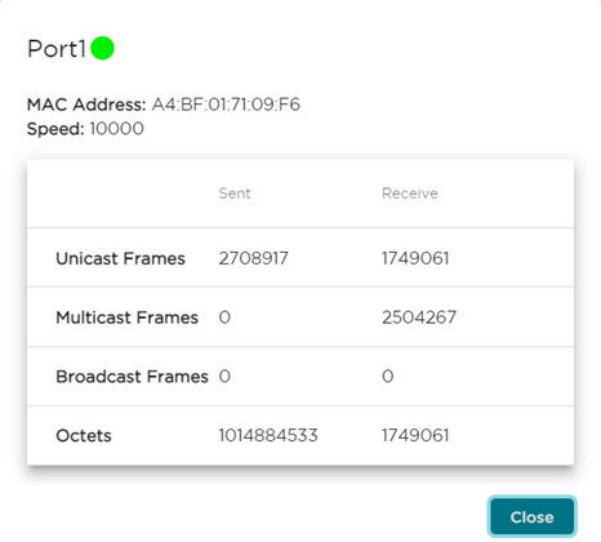 **Figure 13: Port Statistics**

To combine two L2 ports into a LAG port, see Configure LAG Ports on page 59.

ICC Interfaces

The **ICC Interfaces** section displays details of the ICC (Inter-Cluster Connection) interface. The ICC interface is a backplane connection between all the members of a cluster and is used for cluster operations, component state, and shared filesystem synchronization. In a multi-node cluster, each node requires an ICC connection.

The following table describes the fields that display for ICC interfaces.

Table 15: ICC Interface Description

ICC Interface Setting	Description
Name	Name of the interface.
IP Address	IP address of the ICC interface.
MAC	MAC address of the ICC interface.

Table 15: ICC Interface Description (continued)

ICC Interface Setting	Description
Certificates	Any certificates that are assigned to the port.
LAG	Displays whether LAG is enabled or disabled on the interface, based on whether the LAG toggle is enabled or disabled.

To combine two ICC ports into an ICC LAG port, see Configure LAG Ports on page 59.

Interfaces

The **Interfaces** section displays details of the currently configured data interfaces. You can use this section to add network topologies. Topologies represent the networks with which the Universal Compute Appliance interacts. The following table describes the fields that display.

Table 16: Interface Settings

Interface Settings	Description
Name	Name of the interface.
VLAN ID	VLANS that are assigned to the interface.
Tagged	Displays whether the VLAN is tagged (checkmark) or not (no checkmark).
Port	The port connection for the interface.
IP Address	The IP address of the port.
Mode	The mode of the interface. The options are: • Physical - The topology is the native topology of a data plane and it represents the actual Ethernet ports. • Management - The native topology of the [appliance] management port. • Routed - The controller is the routing gateway for the routed topology. • Bridged at Controller - The user traffic is bridged (in the L2 sense) between wireless clients and the core network infrastructure. • Bridged at AP - The user traffic is bridged locally at the AP without being redirected to the controller.
Certificates	Displays certificates that are assigned to the port.

Select **Add New Interface** if you want to configure a new interface. For details, see Add Interface on page 56 .

Static Routes

The **Static Routes** section displays information for currently configured static routes. Use static routes to set the default route of the Universal Compute Appliance so that

device traffic can be forwarded to the default gateway. The following table describes the fields that display.

Table 17: Static Routes Status Fields

Interface Settings	Description
Destination	IP address of the destination Universal Compute Platform.
CIDR	CIDR field is used along with IP address field to find the IP address range.
Gateway	Gateway address of the Universal Compute Platform for any Admin or physical interfaces (B@AC L3 VLAN).
Interface	

Select **Add New Route** if you want to configure a new static route. For details, see Add Static Route on page 60.

Add Interface

Use this procedure to configure a new interface on Universal Compute Platform.

Note
For detailed help with the fields and their configuration options, see Create New Interface Settings on page 57.

1. Go to **Administration** > **System** > **Network Setup**.
2. Under **Interfaces**, select **Add New Interface**.
3. Assign a **Name** and **IP Address** to the interface.
4. Assign any additional settings that you want.
5. Optional. To add a VRRP address on this interface, select **Add** and then configure VRRP settings.
6. Select **Save**.

Create New Interface Settings

The following table provides interface properties for the settings in the **Create New Interface** window.

Table 18: Interface Field Descriptions

Field	Description
Name	Name of the interface.
Mode	Describes how traffic is forwarded on the interface topology. Options are: • Physical - The topology is the native topology of a data plane and it represents the actual Ethernet ports. • Management - The native topology of the Universal Compute Appliance management port. • Routed - The controller is the routing gateway for the routed topology. • Bridged at Controller - The user traffic is bridged (in the L2 sense) between wireless clients and the core network infrastructure. • Bridged at AP - The user traffic is bridged locally at the AP without being redirected to the controller.
VLAN ID	ID for the virtual network.
Tagged	Indicates if the interface tags traffic. When traffic is tagged, the VLAN ID is inserted into the packet header to identify which VLAN the packet belongs to. Tagging can identify the port or interface to send a broadcast message to.
Port	Physical port on the Universal Compute Platform for the interface.
Management Traffic	Enable or disable Management Traffic through this interface. Enabling management provides access to SNMP (v1/v2c, v3), SSH, and HTTPs management interfaces.
MTU	Maximum Transmission Unit (MTU). Standard is 1500 bytes. Fixed value.
IP Address	For an Admin topology, the Layer 3 check box is selected automatically. The IP address is mandatory for a Physical topology. This allows for IP Interface and subnet configuration together with other networking services.
CIDR	CIDR field is used along with IP address field to find the IP address range.

Table 18: Interface Field Descriptions (continued)

Field	Description
FQDN	Fully-Qualified Domain Name
VRRP	This field lets you assign VRRP virtual IP addresses to this interface. For Managed Orchestration clusters, VRRP addresses support load balancing and high availability for cluster services. For Self-Orchestration standalone clusters, VRRP can be used to create an IP alias with access to the application management interface. To add a VRRP address, select **Add**, and configure the following: **Address** Enter the VRRP IP address. Note the following requirements: • You can add up to ten VRRP addresses per data port. However, for ICC ports, you can add a single VRRP address only. To delete an address, select the adjacent icon. • VRRP IPs must be on the same segment as the interface IPs. • VRRP IP addresses must not overlap with any other address in the segment, including addresses assigned by the application. • With clustered deployments, make sure to record the IP address relationship between the cluster's direct interfaces (ICC, Service/Data ports), VRRP, and external access. **Comments** An optional text comment for each VRRP address. Use the comment to describe how that VRRP address is used. **Priority** A numeric value (between 1-254) that determines mastery of the state of exchanges across cluster nodes. The value that you enter applies to all VRRP addresses on this port. As a best practice, note the following: • On a single node, assign the same priority to all Service Set and ICC VRRP addresses. However, assign a different value on different cluster nodes. • As a best practice, designate node 1 as highest priority, node 2 as second highest priority, node 3 as lower priority, and so on in descending order. **Router ID** A numeric value (between 1-255) that allows segmentation of a routing domain. The ID applies to all VRRP addresses on this port. • Within a cluster, assign the same Router ID to the same VRRP addresses across the cluster. • We recommend that you use a different Router ID for ICC VRRPs than for Service Set VRRPs. • It is important to separate this VRRP config from other VRRP uses on the same network segment. The assigned value is arbitrary, but the value must not overlap when another VRRP usage is visible in the attached network segment. **Note:** The **Priority** and **Router ID** settings are important in multi-node clusters only. For Self-Orchestration standalone clusters, assign numeric values to each, but the specific values are not important.

LAG Interfaces

Universal Compute Platform supports the IEEE 802.3ad implementation of Dynamic Link Aggregation Group (LAG), with control managed by the Link Aggregation Control Protocol (LACP). When you join two or more ports into a LAG interface, the network bonds the ports and treats them as a single logical port interface. LAG interfaces increase link throughput and provide redundancy in case of a link failure.

Consider the following when configuring LAG:

- Supported port combinations are: ICC1 and ICC2, and any combination of two to four data ports so long as the ports are configured to run at the same speed.
- An ICC port and a data port cannot be combined into the same LAG interface.
- A single port cannot be added into more than one LAG interface.
- The LAG interface inherits VLAN assignments automatically from newly added port members.
- Universal Compute Platform supports LACP-based LAG only. However, engine applications that support LAG, and which run on Universal Compute Platform, must use static LAG only. The LACP configuration from Universal Compute Platform creates and manages the aggregated link, and the static LAG configuration from the engine runs on that link.
- When deploying LAG, the LACP expiry timeout on Universal Compute Platform is set to **long**. For the switch on the other end of the LAG connection, configure the LACP expiry timeout to **long**.

Configure LAG Ports

Use this optional procedure to configure Link Aggregation Group (LAG) interfaces on Universal Compute Platform.

Note
Make sure to configure LAG on the switch that connects to the LAG ports. Otherwise, the LAG connection fails.

1. Go to **Administration** > **System** > **Network Setup**.
2. To aggregate data ports, under **L2 Ports**, assign data ports to each LAG interface that you want to configure:
   a. For **LAG1**, select each data port from the **Member Ports of LAG** column that you want to add to this LAG.
   b. For **LAG2**, select each data port from the **Member Ports of LAG** column that you want to add to this LAG.
   c. Select **Save** and then select **OK**.

   The LAG interface inherits the VLAN assignments automatically from newly added port members.
3. To aggregate the ICC ports, under **ICC Interfaces**, select **LAG**.
4. Complete the LAG interface configuration for each new LAG interface:
   a. Under **Interfaces**, select the new LAG interface.
   b. Configure the interface settings. For more information, see Create New Interface Settings on page 57.

c. Select **Save**.

Add Static Route

Static Routes define the default route to Universal Compute Platform for legitimate wireless traffic. You must be a system administrator to add a static route.

Note
Static Routes affect the settings for the Default Gateway IP address under **Host Attributes**. Adding a default static route (0.0.0.0/0) changes the Default Gateway IP address.

To add a static route, take the following steps:

1. Go to **Administration** > **System** > **Network Setup**.
2. Under Static Routes select **Add New Route**.

   The **New Static Route** dialog displays.
3. Configure the Static Route Properties. For detailed field descriptions, see Static Route Properties on page 60.
4. Select **Save**.

Static Route Properties

Details about Static Route Properties.

Table 19: Static Route Parameters

Field	Description
Destination	IP address of the destination Universal Compute Platform.
CIDR	CIDR field is used along with IP address field to find the IP address range.
Gateway	Gateway address of the Universal Compute Platform for any Admin or physical interfaces (B@AC L3 VLAN).

Certificates

To view certificate information for specific interfaces, go to **Administration** > **System** > **Network Setup**, select the interface for which you want to view certificate information, and then select **Certificates**.

From this window, the certificate status fields provide information for the certificate that is currently applied to the interface.

Table 20: Certificate Status Fields

Field	Description
Name	The name of the interface.
Expiry Date	Expiration date of the certificate.

Table 20: Certificate Status Fields (continued)

Field	Description
CA Cert	Indicates whether the certificate includes a root certificate chain: • ⚪—A grey dot indicates that the certificate does not include a root certificate chain. • 🟢—A green dot indicates that the certificate includes a root certificate chain.
Name (CN)	The fully qualified domain name for the server that uses this certificate.
Org Unit (OU)	The name of the organization unit or division that uses this certificate.
Organization (O)	The legal name of the organization or company that uses this certificate.

By default, Universal Compute Platform comes with a factory-installed self-signed certificate that applies to all interfaces, and which uses a certificate common name of *Network Services Engine*. The default certificate is used by the user interface HTTP Server to terminate the HTTPS browser requests served on port 5825.

However, you also have an option to upload a public CA-signed certificate to an interface that replaces the self-signed certificate on that interface only.

Go to Update Certificates on page 61 if you want to upload a public CA certificate for use on an interface.

Update Certificates

Use the following optional procedure to install a public CA certificate on an interface or to revert an interface to use the default self-signed certificate.

1. Go to **Administration** > **System** > **Network Setup**.
2. Select the interface for which you want to update certificates.
3. Select **Certificates** to view certificate information for that interface.

4. Select one of the following four certificate options and complete the required substeps for that option:

Table 21: Certificate Update Options

Field	Description
Replace/Install Interface Certificate	Select this option if you want to generate a CA certificate for this interface using Universal Compute Platform's CSR request form. Complete the following substeps: a. Go to Generate CSR on page 63 and create a certificate request that you can present to a CA in exchange for a CA certificate. b. Once you've obtained the certificate, select the first **Choose File** and then select the certificate (.cer file). c. Optional. If you have a root certificate chain, select the second **Choose File** and select the certificate chain (.PEM file).
Replace/Install Interface Certificate and key from a single file	Select this option if you want to install a certificate file that you generated on an external CA site. Complete the following substeps: a. Select the first **Choose file** and select the PKCS#12 certificate (.pfx file). b. Enter the **Private Key** for the certificate. c. Optional. If you have a root certificate chain, select the second **Choose File** and then select the certificate chain (.PEM file).
Replace/Install Interface Certificate and key from separate files	Select this option to upload certificates where the certificate, key, and CA root chain are in different files. For example, this may occur if you generated the files in Linux with an OpenSSL connection: Complete the following substeps: a. Select the first **Choose File** and select the certificate file (.cer file). b. Select the second **Choose File** and select the private key (.key file). c. Optional. If you have a root certificate chain, select the third **Choose File** and select the public CA root chain (.PEM file).
Reset Interface to the factory default certificate and key	Select this option if you want to reset this interface to use the factory-default self-signed certificate. When you select this option, any certificates that you installed on this interface get removed.

5. Select **Save**.

Note
A server restart is triggered after cerificates are applied on the following topologies:

- When applied or reset on the Admin topology.
- When applied or reset on System topologies where Management Traffic is enabled.

Generate CSR

Use this procedure to complete a Certificate Signing Request (CSR) that you can use to generate a CA certificate that you can install on a Universal Compute Platform interface. You can issue a CSR from the **Certificates** window.

1. Go to **Administration** > **System** > **Network Setup**.
2. Select an interface and select **Certificates**.
3. In the **Certificates** window, select **Generate CSR**.
4. Complete the fields in the **Generate Certificate Signing Request** window.

Table 22: Generate Certificate Signing Request Fields

Field	Description
Country Name	Enter the two-letter ISO abbreviation for the country.
State or Province	Enter the name of the state or province.
Locality Name	Enter the locaility name (for example, the city or town).
Organization Name	Enter the name of the organization.
Organizational Unit	Enter the name of the unit within the organization.
Common Name	Enter the certificate common name (e.g., FQDN or IP address).
Email Address	Enter an email address for notification purposes.
Key Size	Select the key size in bits. Supported sizes are 1024 or 2048.

5. Select **Save**.

Generate and download the CSR request that you can present to a public certificate authority. Once you've received a CA certificate back from the certificate authority, return to Update Certificates on page 61 to install the certificate on an interface.

Network Time

System administrators can configure network time and add NTP servers. Go to **Administration** > **System** > **Network Time**.

Table 23: Network Time Settings

Field	Descriptions
System Time	Displays the current system date and time.
Configured Time Zone	Displays the current time zone setting.
Set New Time Zone	To set a new time zone, from the drop-down, select a new time zone and then selevt **Save** to manually change the system date and time.
NTP	Select this option to configure servers for Network Time Protocol (NTP). NTP is an Internet Standard Protocol that assures accurate synchronization to the millisecond of computer clock times in a network of computers.

Table 23: Network Time Settings (continued)

Field	Descriptions
NTP Reachable	An icon indicates if the NTP server is reachable: For example, • A green icon indicates that the server is reachable. • A red icon indicates that the server is not reachable. Check your NTP server settings.
NTP Server 1	Enter the IPv4 address or fully qualified domain name of the primary NTP server.
NTP Server 2	Enter the IPv4 address or fully qualified domain name of the secondary NTP server.

Select **Save** to save updates to the settings.

Settings

Cloud Visibility

If your deployment is onboarded to ExtremeCloud IQ, you can view the cloud address from **Administration** > **System** > **Settings**. This page populates automatically when you onboard the cluster to ExtremeCloud IQ. For example, the URL may look like:

`<RDC name>-cw.extremecloudiq.com` where:

- `<RDC name>` is your Regional Data Center (RDC) information available under **About ExtremeCloud IQ**.
- `-cw` indicates a Universal Compute Platform appliance.
- `.extremecloudiq.com` is the ExtremeCloud IQ host address.

Add Web Proxy Server

For enhanced data security, you can add a web proxy server. A proxy server is an additional server in a client-server deployment that provides additional data security boundaries, protecting users from malicious activity on the internet.

1. Select the navigation menu and select **Administration** > **System** > **Settings**.
2. Select the **Web Proxy** tab.
3. Enter the **IP Address** of the proxy server along with the server **Port** to which you should connect.
4. If the proxy server requires authentication, select **Authentication** and enter the **Username** and **Password** for an account that has access to the proxy server.
5. Select **Save**.

Upgrade the Universal Compute Platform

You can access options that let you upgrade Universal Compute Platform software at **Administration** > **System** > **Software Upgrade**.

The user interface displays information under the following tabs:

- **Image Management**—This tab lets you view the software images that have been uploaded to this appliance. You can select an image and select one of the following icons to complete an action using that image:
  - (Copy to Nodes)—Copy this image to other nodes in the cluster.
  - (Delete)—Delete the image from the appliance.
  - (Upgrade)—Start an upgrade using this image.
  - (Refresh)—Refresh the screen.
- **Upload**—This tab lets you upload new images for Universal Compute Platform.
- **Schedule**—This tab lets you configure an upgrade schedule.
- **Kubernetes**—This tab displays a list of nodes with the current Pod version and Kubernetes version for each node. All nodes should be running the same Pod and Kubernetes version.
- **Logs**—This tab contains logs with information about upgrade history, upgrade details, and restore history.

Go to Upgrade Universal Compute Platform Task Flow on page 65 to initiate the upgrade process.

Upgrade Universal Compute Platform Task Flow

To upgrade Universal Compute Platform, complete the tasks in the following task flow.

Table 24: Upgrade Universal Compute Platform Task Flow

	Procedure	Description
1	Upload Software Image on page 66	Upload the new software image to a Universal Compute Platform cluster node.
2	Copy Image to All Nodes on page 66	(Clustered deployments only). Copy the uploaded image to other cluster nodes.
3	Upgrade the cluster nodes using either of these procedures: • Upgrade Nodes on page 67 • Schedule an Upgrade on page 68	You can initiate an immediate (on-demand) upgrade of cluster nodes or schedule the upgrade for the future. **Note:** With either option, you must upgrade the nodes one node at a time.
4	Kubernetes Upgrade on page 68	Check that your Kubernetes version is current across the cluster and upgrade the version if it's not.

Note

- For information on past upgrades, see Upgrade Logs on page 69.
- For information on how to upgrade container applications, see Engine Upgrades on page 35.

Upload Software Image

Use this procedure to upload a new image file to a Universal Compute Platform node. You must upload the image before you can use the image to complete an upgrade.

Note
The software image must be accessible from your local computer.

1. Go to **Administration** > **System** > **Software Upgrade**.
2. Select **Upload**.
3. For **Image Type**, select **Upgrade** or **Backup**, depending on the type of image.
4. For the **Destination**, select **Local**.
5. Select the **Upload Method** (HTTP, FTP, or SCP).
6. Complete one of the following actions according to the selected upload method
   - For **HTTP** uploads, complete one of the following options to upload the file:
     - Select and drag the image file to the Universal Compute Platform desktop.
     - Select the ↥ (Choose Upgrade File) icon and then browse to the image file and select it.
   - For **FTP** or **SCP** uploads, complete the additional server fields that display according to the below requirements and then select **Upload Image**:
     - **Server IP**—Enter the IP address of the server where the image is stored.
     - **Username**—Enter a username for an account that has access to the server.
     - **Password**—Enter the password for the preceding user account.
     - **Directory**—Enter the directory where the software image is stored.
     - **Filename**—Enter the filename of the software image file.

   The image file uploads to Universal Compute Platform.

What to do Next

For clustered deployments, Copy Image to All Nodes on page 66.

Otherwise, upgrade this node using one of the following procedures:

- Upgrade Nodes on page 67
- Schedule an Upgrade on page 68

Copy Image to All Nodes

For clustered deployments, use this procedure to copy the software image from one cluster node to other nodes in the cluster.

Note
- The image must be uploaded already to the source node for the Copy.
- The Copy feature applies to Universal Compute Platform software installation images only.

1. Log in to the Universal Compute Platform node where the image is uploaded.
2. Go to **Administration** > **System** > **Software Upgrade**.

3. Select **Image Management**.
4. Select the image that you want to copy and then select the (Copy to Nodes) icon.
5. In the **Copy image to nodes** popup, set the following fields:
   - **Image**—Make sure that the correct file is selected.
   - **Copy Image to**—Select each destination node for the copy. You can select multiple nodes.

   Note
   The destination nodes must be running version 5.07.01 or later.

6. Select **Copy**.

   The software image copies to the selected nodes.

What to do Next

Upgrade the cluster nodes using one of these procedures:

- Upgrade Nodes on page 67
- Schedule an Upgrade on page 68

Upgrade Nodes

Use this procedure to initiate an on demand upgrade of Universal Compute Platform nodes. You can upgrade each node in the cluster from a single node.

Note
- You cannot upgrade more than one node in the cluster at the same time.
- The software installation image must be uploaded already to the local node (the node on which you've logged in). If you're upgrading a different node than the local node, the software installation image must have been uploaded to that node as well.

1. From any cluster node, go to **Administration** > **System** > **Software Upgrade**.
2. Select **Image Management**.
3. Select the image that you want to use for the upgrade and then select the (Upgrade) icon.
4. Set the following fields in the **Software Upgrade** popup:
   - **Image**—Make sure that the image that you want to use is selected.
   - **Backup System Image to**—Select **Local**.
   - **Upgrade**—Select **Now**.
   - **Node**—Select the node that you want to upgrade.
5. Select **Upgrade**.

   The upgrade process begins.

After the upgrade finishes, restart the procedure and select a different node for upgrade.

Schedule an Upgrade

Configure an upgrade schedule for the local Universal Compute Platform image.

Note

- You can schedule an upgrade for the local node only. For clusters, you must configure an upgrade schedule for each node separately from that node.
- You can upgrade only one node in the cluster at a given time.
- The software image must be uploaded already to the local node.

1. Go to **Administration** > **System** > **Software Upgrade**.
2. Select **Schedule**.
3. Select the **Image** that you want to use for the upgrade.
4. From the **Backup System to** drop-down, select the destination for the backup file:
   - **Local**—Backup file is saved locally.
   - **Flash**—Backup file is saved to flash.
   - **No Backup**—No backup file is created.
5. Assign the following details to the backup:
   - **Backup Filename** that you want to assign.
   - **Timezone** of the appliance.
   - **Time** of the upgrade in 24 hour format HH-MM.
   - **Date** of the upgrade in MM/DD format.

   

   Note
   When you supply a Date and Time that is in the past, the schedule is set for the following year at the specified date and time.

6. Select **Schedule**.

Repeat this procedure on the other cluster nodes to schedule upgrades for those nodes.

Kubernetes Upgrade

Use the **Kubernetes** tab to manage Kubernetes versions. The tab displays a list of nodes, Kubernetes versions, and whether an upgrade is needed. All nodes must be running the same Kubernetes version.

You can upgrade Kubernetes while working either online or offline:

- For online upgrades—The updated Kubernetes version gets downloaded from Docker Hub automatically when you upgrade.
- For offline upgrades—You must upload a Kubernetes package manually before you complete the upgrade.

Before You Begin

When the status message indicates that you need to upgrade Kubernetes, before you upgrade, go to **Dashboard** > **Deployment Health** and verify the following:

- Verify the system is in a healthy state.
- Verify all nodes are on the same version of software.
- Verify Deployment Health has all Operational indicators as (Green) NTP is reachable, Kubernetes nodes are healthy, Kubernetes HA cluster is healthy, Load-balancing services are properly configured and accessible.

To upgrade Kubernetes:

1. Go to **Administration** > **System** > **Software Upgrade**.
2. Select the **Kubernetes** tab.
3. Complete the upgrade using an online or offline upgrade method:
   - For online upgrades—Select **UPGRADE** and then select **OK** to begin the upgrade.
   - For offline upgrades—Upload a Kubernetes resource package manually, using either of these methods:
     - Select (**Upload Kubernetes resource package zip file**) and then browse to the image file and select it.
     - Select the image file from your desktop and drag it into the **Upload Kubernetes resource package zip file** area on Universal Compute Platform.

   The upgrade begins automatically after the image upload completes.

> **Note**
> To delete an available image, select the adjacent (Remove Mirror).

Upgrade Logs

The **Logs** tab displays the following information for the appliance:

- Upgrade History
- Upgrade Details
- Restore Details

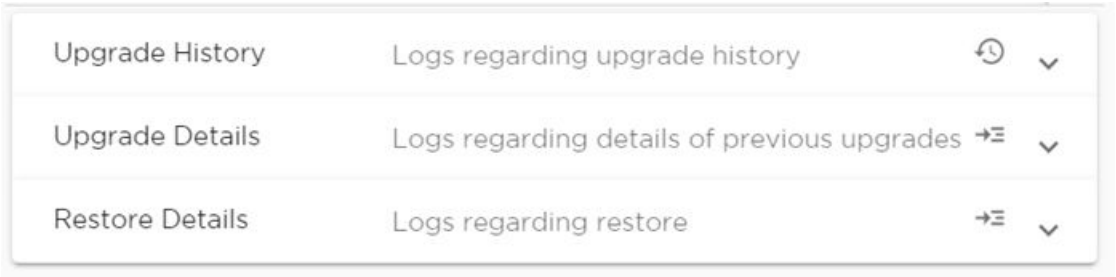

Figure 14: Logs tab

Select to expand each log file.

You can copy text from each log file.

1. Select to expand the log file.

2. Select the log text you want to copy and select .

System Information

Go to **Administration** > **System** > **System Information** to view system and manfacturing information. You can select either of the following tabs:

- **System Information** tab
- **Manufacturing Information** tab

From the **System Information** tab, you can view the following information about your system:

- System Up Time
- CPU Utilization
- Memory Usage—% of free memory.
- Disk Usage:
  - Partitions
  - Total Space
  - Used space
  - Available space
  - Use %
- Flash Usage:
  - Partitions
  - Total Space
  - Used space
  - Available space
  - Use %
- System Temperature—Temperatures for various system components, including:
  - System Boards
  - Front Panel Boards
  - I/O Module
  - System Internal Expansion Boards
  - Drive Backplane
  - Power Supply
  - Processors
  - Memory Modules
- Fan Speed—Revolutions per minute (RPMs) for various system fans.
- Power Supply—Information for each power supply source, including:
  - Model
  - Serial Number
  - Revision
- Port Interface status:
  - Interface state

- Speed of interface

Note
Some fields do not display for all appliance models. For example, Power Supply information does not display on the 1130C.

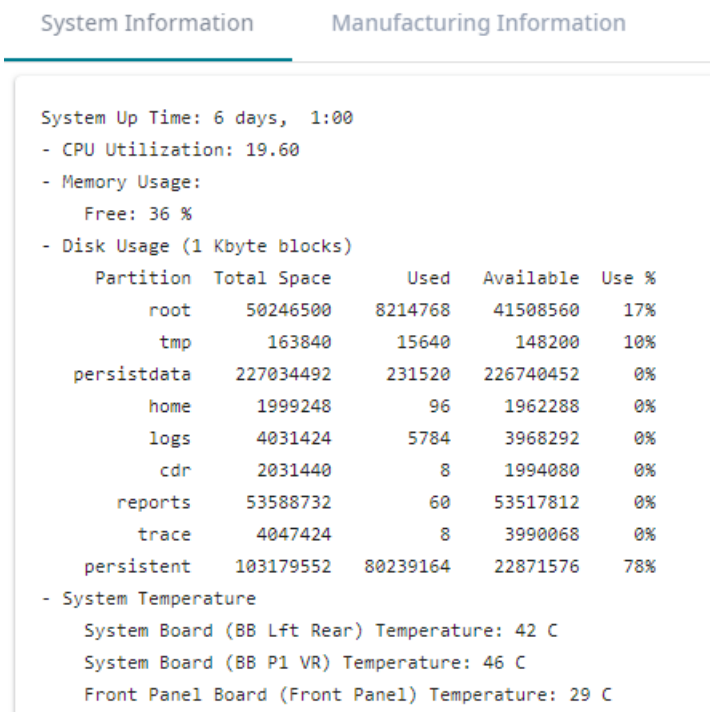

Figure 15: Example System Information

From the **Manufacturing Information** tab, you can view the following information:

- Manufacturing ID (Serial Number)
- BIOS Version
- Hardware Revision
- VF MAC base
- SMX Version
- Software Version
- Model
- CPU Type
- CPU Frequency
- Number of CPUs
- Total Memory
- MAC addresses for various system interfaces
- Locking ID

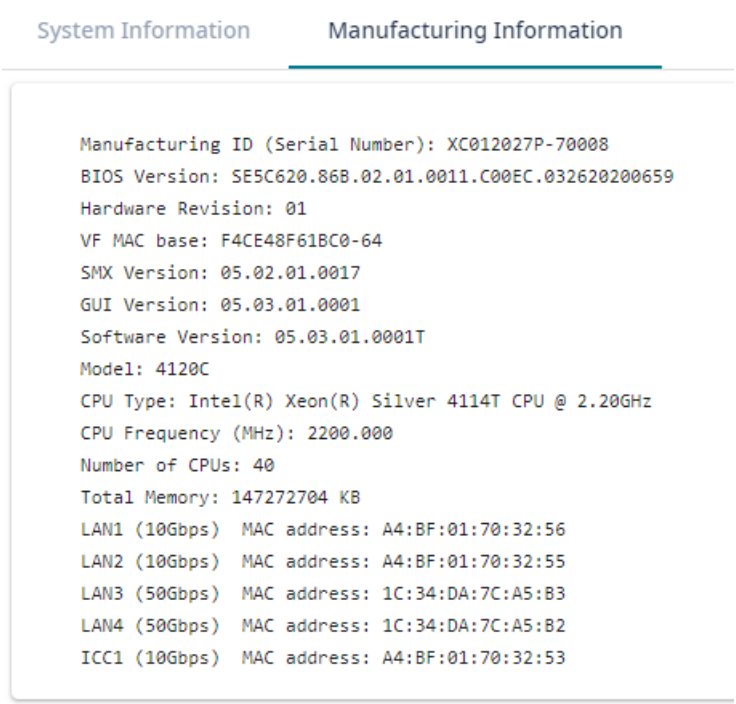

Figure 16: Example Manufacturing Information

Utilities

Universal Compute Platform provides a remote console to a node controller. Use the remote console to open a live SSH console session.

To open a remote console, go to **Administration** > **System** > **Utilities**.

Note

A live console is available from each engine instance for diagnostics and troubleshooting. To open a live console and connect to a container or virtual machine instance (VMI), from the engine **Console** tab, select **Attach**.

Select each engine instance to display engine settings and the **Console** tab for that instance.

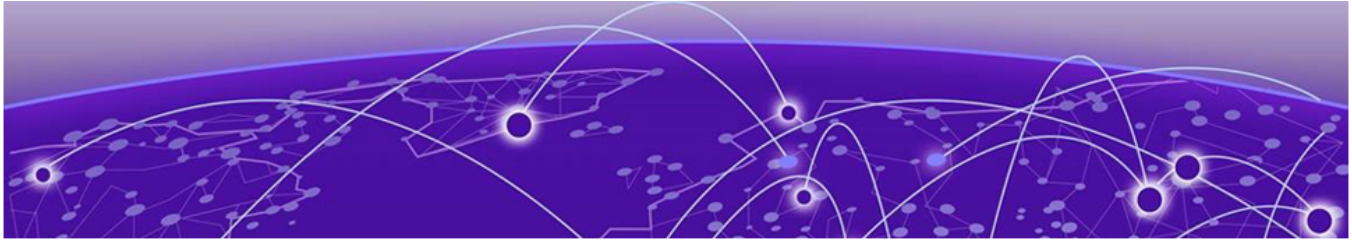

Index

T

U

V

W