



ExtremeCloud Edge v5.14.01 Self-Orchestration Deployment Guide

Configuration and Management for Universal Compute Platform

9041053-00 Rev. AA
June 2026



Copyright © 2026 Extreme Networks, Inc. All rights reserved.

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.....	v
Preface.....	6
Text Conventions.....	6
Documentation and Training.....	7
Open Source Declarations.....	8
Training.....	8
Help and Support.....	8
Subscribe to Product Announcements.....	9
Send Feedback.....	9
Introduction.....	10
Self-Orchestration.....	10
Supported Appliances for Self-Orchestration.....	12
Kubernetes.....	13
Network Architecture.....	14
High Availability.....	16
Requirements.....	17
Firewall Requirements.....	17
Firewall Requirements for Online Installs and Upgrades.....	17
Firewall Requirements for Cloud Connected Deployments.....	18
Reserved IP Addressing.....	18
VRRP Configuration (Optional).....	18
Configure an Appliance.....	20
Connect to the Appliance through the Console Port.....	21
Connect to the Appliance through the ICC1 Port.....	21
Run the Basic Configuration Wizard.....	21
Upgrade the Appliance Universal Compute Platform.....	25
Validate the Network Address Configuration.....	27
Add a Port.....	29
Install Kubernetes (Offline Installation).....	31
Configure the Stand-Alone Cluster Settings.....	32
Engine Application Installation.....	34
Download the Docker Application Image.....	35
Upload Application Image to Appliance.....	35
Install Engine Application.....	35
Deploy Application.....	35
Configure Interface Settings for Engine Application.....	36
Engine Application Settings.....	36
Engine Upgrades.....	37
Upgrade an Application (Self-Orchestrated).....	38

Onboard Cluster to ExtremeCloud IQ.....	39
Onboarding a Cluster to ExtremeCloud IQ	39
Cloud Visibility.....	40
Index.....	41



Abstract

The ExtremeCloud Edge v5.14.01 Self-Orchestration Deployment Guide provides in-depth technical guidance for deploying the Universal Compute Platform with self-orchestration capabilities. The v5.14.01 introduces support for offline installations and upgrades, support for Third-Party Management Engine on 1130C, and updates to the VRRP configuration for cases where you need to deploy multiple VRRP addresses on a single data interface. Architecture details for Self-Orchestration focus on the orchestration layer, automated provisioning pipelines, and integration with network and cloud infrastructures. Key deployment processes include configuring system prerequisites, establishing deployment pipelines, and executing end-to-end orchestration workflows. Advanced troubleshooting techniques, optimization strategies, and best practices are provided to resolve common deployment challenges and ensure optimal performance. This resource is designed for system architects and IT professionals overseeing large-scale, automated Universal Compute Platform implementations.



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 <i>enter</i> and <i>type</i>	When you see the word <i>enter</i> 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 <i>type</i> .
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
<i>Words in italicized type</i>	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.
<i>italic text</i>	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, <i>member [member...]</i> .
\	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.



Introduction

[Self-Orchestration](#) on page 10

[Supported Appliances for Self-Orchestration](#) on page 12

[Kubernetes](#) on page 13

[Network Architecture](#) on page 14

[High Availability](#) on page 16

This guide provides the steps needed to bring a stand-alone Universal Compute Platform online in ExtremeCloud Edge - Self-Orchestration mode.

Universal Compute Platform leverages Kubernetes and Docker to deploy and manage the delivery of applications to the customer premises, providing the computing power, storage, high availability, and load-balancing for the system.



Note

The internal Kubernetes engine requires the reservation of two (2x) /16 or smaller subnets. Ensure that this IP address range does not conflict with any routable address space within the organization

Self-Orchestration

Deploy Universal Compute Platform in ExtremeCloud Edge - Self-Orchestraton (standalone mode) with a set of Universal Container applications.

This deployment scenario includes the following application requirements for a stand-alone configuration:

- All IP addresses must be unique within their corresponding segments.
- The use of the Virtual IP Address (VIP) is optional, but for some applications it provides a convenient way to expose services, from an instance of the UI (port 443) externally. Assign a VIP to each instance of an engine.
- Pod Network configuration settings — Pods are 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.
 - Pod Network IP Address and CIDR

- Service Network IP Address and CIDR

**Note**

CIDR (*Classless Inter-Domain Routing*) is a method for allocating IP addresses and for IP routing.

**Note**

It's mandatory to configure an IP address for the ICC1 port, even in a Standalone deployment. However, connectivity for the ICC1 port is not required.

Related Links

[Reserved IP Addressing](#) on page 18

[VRRP Configuration \(Optional\)](#) on page 18

Supported Appliances for Self-Orchestration

ExtremeCloud Edge - Self-Orchestration deployments of Universal Compute Platform support the following hardware appliances. Depending on the hardware, you may be able to install more than one instance of an application on a node.



Note

Support is for a single application type per node. Application mixing on a single appliance is not supported.

Table 4: Supported Hardware for ExtremeCloud Edge - Self-Orchestration

Hardware Appliance	Details
1130C	<p>Ports:</p> <ul style="list-style-type: none"> • 2 x 1 Gbps ICC Ports/RJ45 • 4 x 1 Gbps Data 1-4/RJ45 <p>Self-Orchestration deployment application capacity:</p> <ul style="list-style-type: none"> • Tunnel Concentrator—One instance per node • ExtremeCloud IQ Controller (CE1000)—One instance per node • Third-Party Management Engine—One instance per node <p>For additional server specifications, along with hardware installation information, see Universal Compute Platform Appliance 1130C Installation Guide.</p>
2130C	<p>Ports:</p> <ul style="list-style-type: none"> • 2 x 1/10 Gbps ICC Ports/RJ45 • 2 x 1/10 Gbps Data 1-2/RJ45 • 2 x 10/25 Gbps Data Ports 3-4/SFP28 <p>Self-Orchestration deployment application capacity:</p> <ul style="list-style-type: none"> • Tunnel Concentrator—One instance per node • ExtremeCloud IQ Controller (CE2000)—One instance per node • ExtremeCloud IQ - Site Engine—One instance per node • ExtremeControl—One instance per node • ExtremeAnalytics—One instance per node <p>For additional server specifications, along with hardware installation information, see Universal Compute Platform Appliance 2130C Installation Guide.</p>

Table 4: Supported Hardware for ExtremeCloud Edge - Self-Orchestration (continued)

Hardware Appliance	Details
3150C	Ports: <ul style="list-style-type: none"> • 2 x 1/10 Gbps ICC Ports/RJ45 • 2 x 10/25 Gbps Data Ports 1-2/SFP28 • 2 x 10/25/50/100 Gbps Data Ports 3-4/QSFP28 Self-Orchestration deployment application capacity: <ul style="list-style-type: none"> • Tunnel Concentrator—One instance per node • ExtremeCloud IQ Controller (CE3000)—One instance per node • ExtremeWireless WiNG Controller (CX9000)—One instance per node For additional server specifications, along with hardware installation information, see Universal Compute Platform Appliance 3150C Installation Guide
4120C/4120C-1	Ports: <ul style="list-style-type: none"> • 2 x 1/10 Gbps ICC Ports/RJ45 • 2 x 1/10 Gbps Data 1-2/RJ45 • 2 x 1/10/25/40/50 Gbps Data 3-4/QSFP Self-Orchestration deployment application capacity: <ul style="list-style-type: none"> • Tunnel Concentrator—Up to three instances per node. • ExtremeWireless WiNG Controller (CX9000)—One instance per node For additional server specifications, along with hardware install information, see Universal Compute Platform Appliance 4120C Installation Guide .

Appliance Installation and Specifications Documentation

For appliance installation and technical specifications information, see the [Universal Compute Platform Appliance Installation Guide](#) for your specific hardware appliance model. These guides contain the following information:

- Verifying the Contents Packaged with the Appliance
- Connecting a Power Supply
- Mounting the Appliance
- Installing Transceivers and Transceiver Cables
- Technical Specifications

Kubernetes

Universal Compute Platform is built on Kubernetes middleware. Kubernetes provides a unifying structure for application delivery and provides integrated management of application state along with clustering capabilities.

Kubernetes can be installed while working online or offline:

- For online installations, Kubernetes components are downloaded and installed during the cluster configuration stage. After you select the cluster type and initialize the cluster (a **Standalone** cluster for Self-Orchestration deployments), the appliance

connects to Docker Hub automatically to download and install the additional Kubernetes components based on your installation requirements. For firewall information for online installations, see [Firewall Requirements](#) on page 17.

- For offline installations, you must upload and install a Kubernetes resource package with the Kubernetes image before you create the cluster. Once Kubernetes is installed, you can create the standalone cluster.

After the Standalone cluster is created, Kubernetes binds to the ICC IP address (either the physical ICC IP address or VRRP IP, if it's configured). Due this binding, an ICC IP address is required for all Self-Orchestrated deployments, although ICC connectivity is not required.



Note

- Because of the ICC binding, it's recommended to use the data ports for application management rather than the ICC ports.
- Do not change the ICC addressing scheme, the hostname, or the domain name, once they are assigned. If you change the ICC IP address or ICC VRRP address, the Kubernetes binding breaks, and the Kubernetes installation unwinds, effectively wiping out the installation. In this case, the only fix is to reinstall and reconfigure. The user interface prevents modification of these parameters to preserve the integrity of the system. If adjustments are required, the node must be reset.
- Kubernetes requires the reservation of two /16 subnets for use by the Pod and Service Networks (the default ranges are 10.96.0.0 and 10.97.0.0). Make sure that the ranges that you use do not overlap with routing domains.




Figure 1: Kubernetes Middleware Layer

Network Architecture

Each Universal Compute Platform appliance has two ICC ports and four data ports.

Access to the data interface of Universal Compute Platform is abstracted through the use of internal virtual switches where each data port has a dedicated vSwitch (four in total). All external connections to the data interface pass through one of these vSwitches.

Each vSwitch has between eight and 16 virtual functions (VFs) that let you run multiple application instances over the same data port. Each VF + vSwitch combination has a unique MAC address. For each installed application, you must allocate a VF on each data port to that instance. Additional application instances must be assigned to a different VF. For example, if you assign the default VF assignments on all ports, a 4120C with three Tunnel Concentrator instances uses VF01, VF02, and VF03 on each port to manage the three application instances.

When you run multiple application instances on an appliance, the instances share the same port and switch, but the attachment point is different at the PCI level. The installed application instances attach to the vSwitch, but the applications see those attachment points as PCI interfaces that are mapped to a VF, and bound to a physical data port.

**Note**

The 1130C has eight VFs for each vSwitch. All of the other supported appliances have 16 VFs per vSwitch. For example, a 2130C has 70 MAC addresses: two MACs for the ICC ports, four MACs for the data ports, plus 64 VF-related MACs.

Universal Compute Platform OS Access

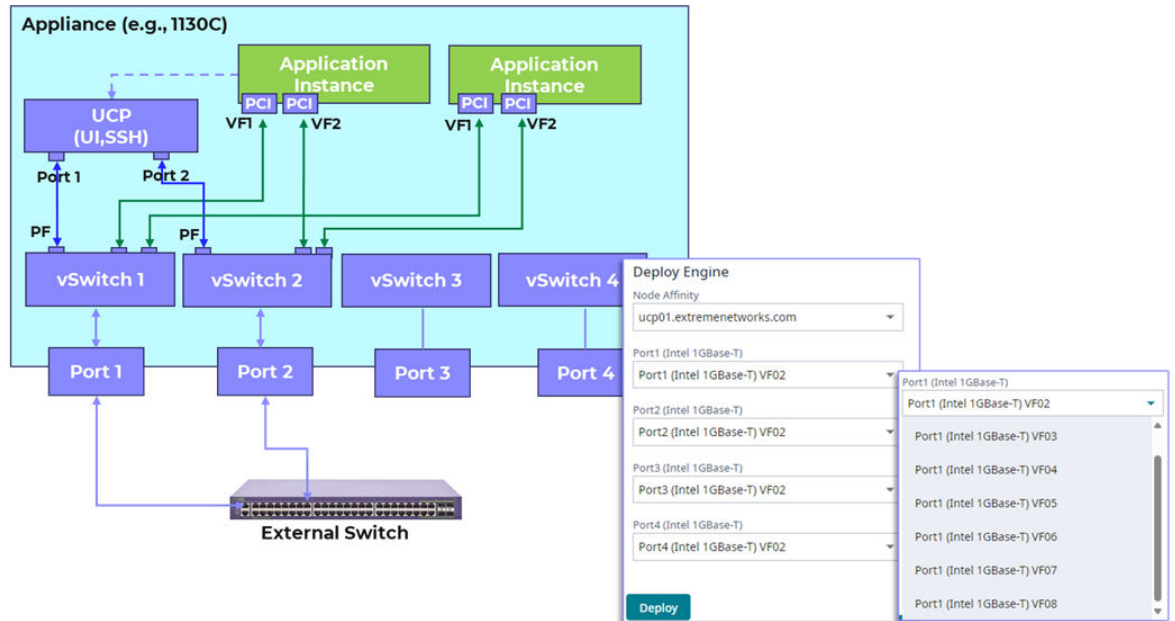
The Universal Compute Platform OS and management application shares access to resources with the installed applications. The Universal Compute Platform attaches to the physical function of the vSwitch, which provides more privileged access to port functions.

The installed applications have an internal hook to the Universal Compute Platform host through an internal NAT interface (10.0.2.2/24) back into the host. Application can use this connection for API REST calls to complete various tasks such as checking interfaces.

Example

The following image illustrates the network connections on an 1130C that has two application instances that use VF01 and VF02 respectively. The Universal Compute Platform OS attaches to the physical function (PF) of vSwitch 1 and vSwitch 2. The user interface callouts display how, for each data port, you can allocate one of eight VFs to that application instance.

Network model



High Availability

Most Self-Orchestrated applications are not designed to use the clustering capability of Kubernetes. For these applications, High Availability is provided by the application, rather than by the Universal Compute Platform.

To configure High Availability for most Self-Orchestration applications, configure separate standalone clusters and then configure HA at the application level. For example, the following image displays a High Availability setup for a CE1000 deployment of ExtremeCloud IQ Controller. In this example, High Availability is provided at the application level.




Figure 2: High Availability for Self-Orchestrated Applications



Requirements

- [Firewall Requirements](#) on page 17
- [Reserved IP Addressing](#) on page 18
- [VRRP Configuration \(Optional\)](#) on page 18

Firewall Requirements

The following sets of firewall requirements exist for Self-Orchestrated deployments. Depending on your deployment type, these requirements may or may not apply to your deployment.

- [Firewall Requirements for Online Installs and Upgrades](#) on page 17
- [Firewall Requirements for Cloud Connected Deployments](#) on page 18

Firewall Requirements for Online Installs and Upgrades

For online installations and upgrades of host firmware, connection to the following domains is required during the installation or upgrade so that Universal Compute Platform can download and install the appropriate packages.

Once the stand-alone cluster is created or upgraded, internet connectivity is not required.



Note

- These requirements do not apply to offline installations and upgrades.
- These requirements do not apply to application engine software that runs on the Self-Orchestration deployment. For example, internet connectivity is not required while installing, upgrading, or running application engines that run on the appliance (unless the specific application engine requires it, or unless you onboarded to the cloud).

Table 5: Required Connection Domains for Online Installs and Upgrades

Domain Name	IP Adresses	Protocol	Port
docker.io	Dynamic IP range	HTTPS	443
gcr.io	Dynamic IP range	HTTPS	443
ghcr.io	Dynamic IP range	HTTPS	443

Table 5: Required Connection Domains for Online Installs and Upgrades (continued)

Domain Name	IP Adresses	Protocol	Port
quay.io	Dynamic IP range	HTTPS	443
registry.k8s.io	Dynamic IP range	HTTPS	443

Firewall Requirements for Cloud Connected Deployments

If you are onboarding to the cloud, connectivity to the following ExtremeCloud IQ services is required.

Table 6: Required Connections for Cloud Deployments

Domain Name	IP Addresses	Protocol	Port	Description
hac.extremecloudiq.com	34.253.190.192 ~ 34.253.190.255	HTTPS	443	Onboarding to ExtremeCloud IQ
<rdc>-inlets.extremecloudiq.com	Dynamic IP range	TCP	8090	Ongoing connection to ExtremeCloud IQ

Reserved IP Addressing

Container orchestration by Kubernetes within the cluster requires reservation of private network segments for each Pod. Plan for network segmentation regardless of your deployment mode.



Note

Review the default IP range values for your pod and service networks in the following table. Use them if they are suitable and do not conflict with the deployed infrastructure network routing definitions. If there is a conflict, adjust the segment IP range as required.

Table 7: IP Address range for network segmentation

Restricted IP Range	Default Value	IP Address /Range
Pod Network IP Range	10.96.0.0/16	<reserved ip>/16
Service Network IP Range	10.97.0.0/16	<reserved ip>/16
Application Network IP Range	10.0.2.0/24	<reserved ip>/24

VRRP operations require visual representation of where the IP addresses are allocated.

VRRP Configuration (Optional)

The Universal Compute Platform relies on Virtual Router Redundancy Protocol (VRRP) to provide IP abstraction to key functionality.

Some applications request that you configure a VRRP address for access to that application's management interface. VRRP can be used to configure an IP alias with pass-through capability for the Universal Compute Platform data interface to that application's management interface.


If you are using VRRP, the following settings must be defined for the VRRP configuration on a data port:

- **Address**—You can assign up to ten VRRP IP addresses for a single data port. Note the following:
 - VRRP IPs must be within the same segment as configured for the interface.
 - VRRP IP addresses must not overlap with any other address in use in the segment, including addresses that are assigned by the application.
- **Priority**—VRRP uses priority settings as a mechanism to arbitrate mastery of the state of exchanges across members of the cluster.
- **RouterID**—This setting allows segmentation of a routing domain, and it is important to separate from any 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 segments.

**Note**

In a stand-alone configuration, configure priority and router ID with a numeric value. However, in a stand-alone configuration, the specific value is not important. These attribute definitions are important in a multiple-node configuration.

For details, see [Add a Port](#) on page 29.



Configure an Appliance

- [Connect to the Appliance through the Console Port](#) on page 21
- [Run the Basic Configuration Wizard](#) on page 21
- [Upgrade the Appliance Universal Compute Platform](#) on page 25
- [Validate the Network Address Configuration](#) on page 27
- [Install Kubernetes \(Offline Installation\)](#) on page 31
- [Configure the Stand-Alone Cluster Settings](#) on page 32

To configure the appliance for an ExtremeCloud Edge - Self-Orchestration deployment, complete the following tasks.

Table 8: Configure an Appliance Task Flow

Step	Procedure	Description
1	Connect to the Appliance through the Console Port on page 21	Connect to the hardware appliance for the initial setup.
2	Run the Basic Configuration Wizard on page 21	Run the wizard to deploy a fully-functioning appliance on a network.
3	Upgrade the Appliance Universal Compute Platform on page 25	Upgrade the Universal Compute Platform appliance software to the latest revision
4	Validate the Network Address Configuration on page 27	Validate network settings and configure additional data plane interfaces, if necessary.
5	Install Kubernetes (Offline Installation) on page 31	For offline installations only, install the Kubernetes image before you create the cluster. Note: For online installations, you can skip this procedure.
6	Configure the Stand-Alone Cluster Settings on page 32	Configure and create the standalone cluster.

What to do Next

Go to the chapter [Engine Application Installation](#) on page 34 to install and deploy the engine application.

Connect to the Appliance through the Console Port

Take the following steps to connect to the appliance through the console port:



Note

Alternatively, you can also use the ICC interface. For details, see [Connect to the Appliance through the ICC1 Port](#) on page 21.

1. Connect the laptop serial port to the console port on the hardware appliance.



Note

If the laptop does not support RS232 interface, then obtain a USB to RS232 converter cable, which supports the following connections:

- RJ45-DB9F cable—for 1130C, 4120C, or 4120C-1
- Null Modem DB9 F-F (Female to Female)—for 2130C, 3150C, or 3160C

2. Using PuTTY, TeraTerm, or another terminal emulator, connect to the serial port connection and configure the following settings for the serial port:

- 115200 baud
- 8 data bits
- 1 stop bit
- Parity none
- Flow control none



Note

The system's default gateway must be pointing to a next hop connection through the service ports.

3. To assign network settings, go to [Run the Basic Configuration Wizard](#) on page 21.

Connect to the Appliance through the ICC1 Port

You can use the ICC1 port to connect to the appliance for initial setup. Note that the ICC1 port has a defined IP address, but you must set your laptop to use the same network segment.

1. Connect a laptop to the appliance management port.
2. Configure the laptop Ethernet port of the laptop with a statically assigned unused IP address in the **192.168.10.0/24** subnet.
3. SSH to the appliance.

192.168.10.1 is the default IP address on the appliance management port).

The Universal Compute Platform logon screen is displayed.

4. To assign network settings, go to [Run the Basic Configuration Wizard](#) on page 21.

Run the Basic Configuration Wizard

Run the Basic Configuration Wizard to assign basic network settings such as IP addresses, VLAN IDs, hostnames, DNS, default gateway, and NTP.

The wizard launches automatically when you log in to the appliance for the first time, or after the appliance has been reset to factory-default settings. You can also run the wizard after the initial setup to update network settings.

**Note**

The wizard prompts you with a series of yes/no, multiple choice, and manual entry questions. The following conventions apply:

- You must press the `Enter` key after each entry to input your entry.
- Displayed settings in `[square brackets]` represent the default value for that prompt, for example `ICC1 IP Address [192.168.10.1]` where 192.168.10.1 is the default IP address for the ICC1 port. To apply the default, just press `Enter`.
- Displayed settings in `(round brackets)` represent a list of options from which you must make a selection, for example `(y|n) [y]` where yes and no are options (and the default is yes). To select no, which is a non-default value, press `n` and then press `Enter`.

**Note**

Settings are configured with the following settings groups. Complete settings for the group and then save those settings before moving on to the next group of settings.

- Admin Password Configuration
- ICC Port Settings
- Data Port Settings
- Host Attribute Settings
- Global Default Gateway Settings
- Time Settings

1. Log in to the appliance:
 - a. Enter the admin credentials to log in to the appliance.

**Note**

The default admin credentials for the first login are as follows. Note that these values are case-sensitive:

- Username: admin
- Password: abc123

As a best practice, we recommend that you change the password immediately after the first login.

- b. Press `Enter` to begin the setup.

2. In the **Admin Password Configuration** section, do the following:
 - a. To change the default admin password (recommended), select `y` and complete the subsequent steps to change the password. Otherwise, select `n` and proceed to step 3.
 - b. Enter a new password that is between 8–24 characters.
 - c. Re-enter the password.
 - d. Press `Enter` to accept the changes.
3. **Configure the ICC Port Settings:**
 - a. Enter an IP address for the ICC1 port.
 - b. Enter a netmask for the ICC1 port.
 - c. If you want to assign VRRP to the ICC ports, at the VRRP prompt, select `y` and then configure the ICC VRRP details. Otherwise, select `n`.
 - d. If you want to assign LAG to the ICC ports, at the LAG prompt, select `y` and then configure LAG details. Otherwise, select `n`.
 - e. Select `y` to accept the ICC configuration.

**Note**

- In a Self-Orchestration deployment, an ICC1 IP address configuration is required, though ICC1 network connectivity is not required. In most cases, the ICC1 interface factory-default settings are sufficient. However, if the defaults conflict with existing routing paths within your organization, modify the configuration to prevent overlap.
- Do not change the ICC addressing configuration once it is assigned and the cluster is created. If you change the ICC IP address or ICC VRRP address, the Kubernetes binding breaks, and the Kubernetes installation unwinds, effectively wiping out the installation. If this occurs, you must reinstall and reconfigure.
- VRRP and LAG are not required in Self-Orchestration deployments.

4. **Configure the Data Port Settings:**
 - a. If you want to edit the data port settings, select `y` and complete the subsequent steps. Otherwise, select `n` to accept the default settings and proceed to step 5.
 - b. Select the port that you want to assign as the main data port.
 - c. Enter the data port IP address.
 - d. Enter the data port netmask.
 - e. Enter the VLAN ID for the data port.
 - f. To enable tagged frames for the data interface VLAN, select `y`. Otherwise, select `n` to use untagged frames.
 - g. To enable management on the data interface, select `y`. Otherwise, select `n`.
 - h. Select `y` to accept all data port settings.
5. **Configure the Host Attribute Settings:**
 - a. Enter the hostname for the appliance in lower case letters.
 - b. Enter the domain name for the appliance.

- c. Enter the IP address of the primary DNS server.
 - d. If you want a secondary DNS server, select `y` and enter the IP address of the secondary DNS server. Otherwise, select `n`.
 - e. Select `y` to accept all host attribute settings.
- 6. Configure the Global Default Gateway Settings:**
- a. Enter the IP address of the default gateway. The address must point to a next hop connection through one of the service ports.
 - b. Select `y` to accept the default gateway settings.
- 7. Configure the Time Settings:**
- a. To update the time zone of the appliance, select `y` and complete the following Region substeps to configure the time zone. Otherwise, select `n` and go to 7b.
 - i. For Region, select the number for the desired continental region.
 - ii. For Region, select the number for the desired city region.

The configured time zone displays, for example `America/New_York`.
 - b. Enter the IP address or fully qualified domain name (FQDN) of the primary NTP server.
 - c. If you want to add a second NTP server, select `y` and enter the IP address or FQDN of the secondary NTP. Otherwise, select `n`.
 - d. Select `y` to accept the updated time settings.

**Note**

Make sure that the NTP settings are correct before you accept settings. Several system functions depend on an accurate timestamp.

The **Controller Post Installation Configuration** screen displays with a set of options that match the settings groups in the wizard.

Table 9: Controller Post Installation Configuration

Menu Option	Command
Admin Password Configuration	1
Change ICC Port Settings	2
Change Data Port Settings	3
Change Host Attribute Settings	4
Change Global Default Gateway	5
Change Time Settings	6
Apply Settings and Exit	A
Exit Without Applying	E

8. To apply settings and exit, select A.



Note

- If you want to reconfigure any of the previous settings groups or exit without applying the configuration changes, enter the corresponding numbers or characters, as displayed on the **Controller Post Installation Configuration** screen (see [Table 9](#)).
- If you reconfigure any screen, you must also reconfigure all subsequent settings. For example, if you reconfigure the Admin Password, you will have to reconfigure all the subsequent configuration wizard settings.

What to do Next

After you run the configuration wizard, use the `ping` command to test connectivity to external services.

1. To test connectivity to external services, ping the IP address of the external server. For example, to test connectivity to DNS, ping the DNS server.
2. Ping the cluster IP address to test connectivity.

```
Extreme Universal Compute Platform
Copyright Extreme Networks Inc. 2022

-----+
c2-xca4.pinewoods.tor.lab.local# ping
Usage: ping [source-interface (name <name>) | (number <id>)] <ip address>
c2-xca4.pinewoods.tor.lab.local# ping 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_seq=1 ttl=112 time=2.82 ms
64 bytes from 8.8.8.8: icmp_seq=2 ttl=112 time=2.05 ms
64 bytes from 8.8.8.8: icmp_seq=3 ttl=112 time=2.01 ms

--- 8.8.8.8 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2003ms
rtt min/avg/max/mdev = 2.008/2.293/2.818/0.371 ms
c2-xca4.pinewoods.tor.lab.local#
```

Figure 3: Example ping command

Upgrade the Appliance Universal Compute Platform

A best practice is to upgrade the appliance to the latest revision. Take the following steps to upgrade the Universal Compute Platform for the appliance:

1. Download the Universal Compute Platform image file from the [Extreme Networks Support Portal](#). The image file extension is one of the following file types:
 - .asx (for 1130C)
 - .gbx (for 2130C)
 - .ygy (for 3150C)
 - .rcx (for 4120C)

2. Log in to the appliance Admin user interface: `https://node_ip:5825`
3. Go to **Administration > System > Software Upgrade > Upload**.

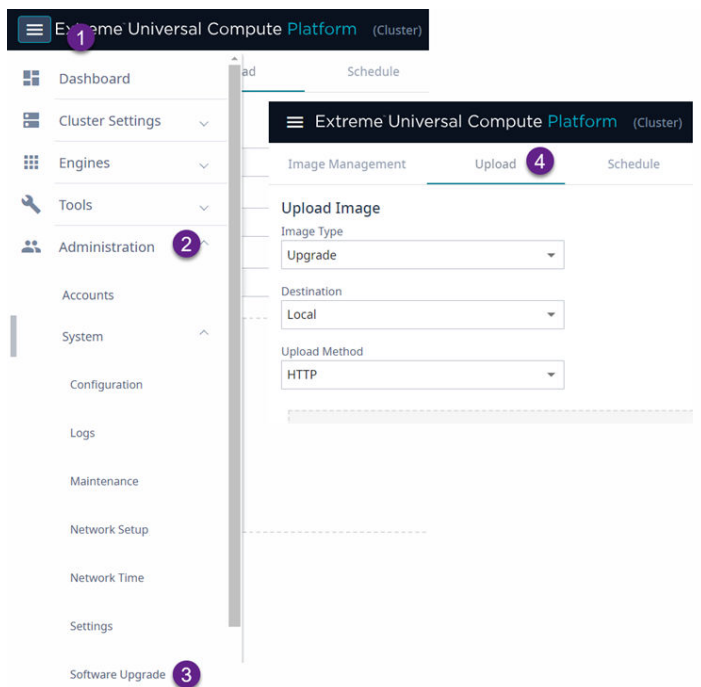


Figure 4: Navigate to Universal Compute Platform Image Upload

4. Specify the Image upgrade settings:
 - Image Type
 - Destination
 - Upload Method. The available upload methods are HTTP, FTP, and SCP; HTTP is recommended.

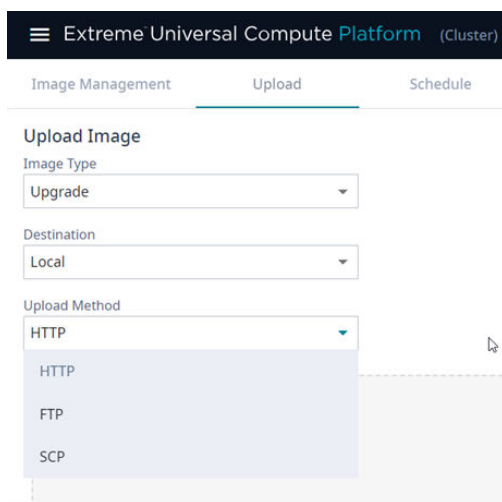


Figure 5: Upload Image Settings

- Upload the desired revision of Universal Compute Platform.
Select the **Choose Upgrade file pane** and navigate to the upgrade image or drag and drop the file on the upgrade pane.

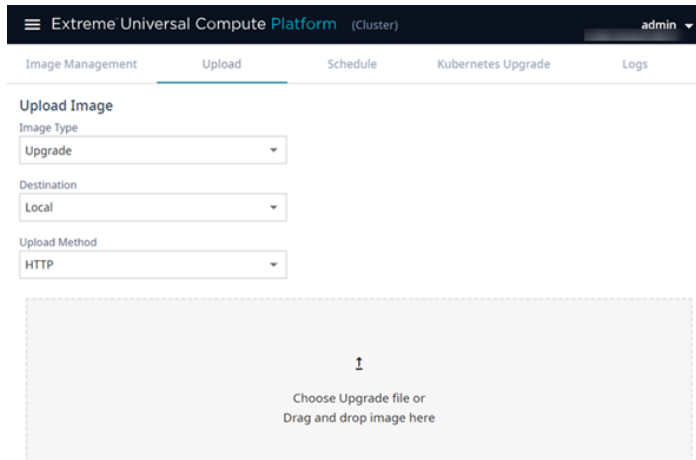


Figure 6: Select the upgrade image



Note

The upgrade may take up to five minutes.

- From the **Image Management** Tab, select the Upgrade image, and click **Upgrade**.

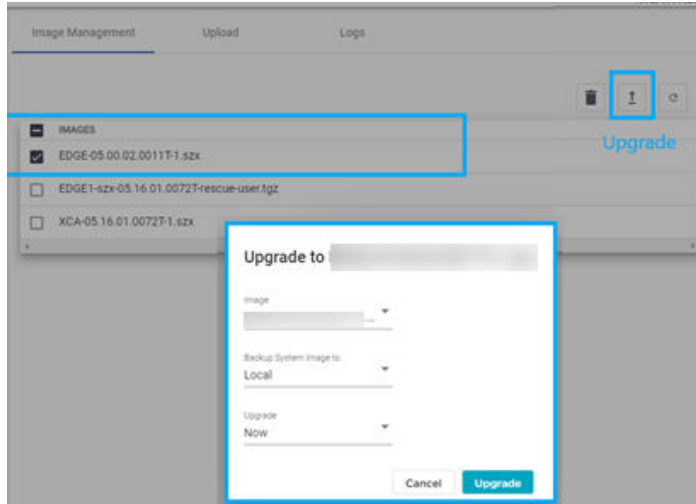


Figure 7: Upgrade the selected image

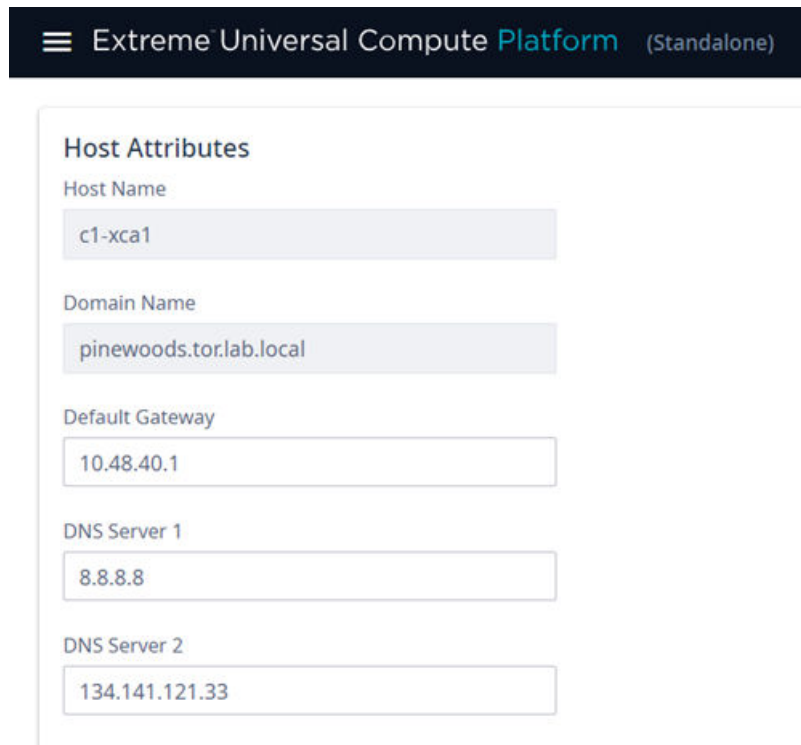
When the cluster is upgraded to the latest revision, proceed to [Validate the Network Address Configuration](#) on page 27.

Validate the Network Address Configuration

Validate the IP addresses that you configured previously through the Configuration Wizard.

To access the network settings in the user interface:

1. Go to **Administration > System > Network Setup**.
2. Verify the host attributes.



The screenshot shows the 'Host Attributes' configuration page in the Extreme Universal Compute Platform (Standalone) user interface. The page is titled 'Host Attributes' and contains several input fields for network configuration:

- Host Name:** c1-xca1
- Domain Name:** pinewoods.tor.lab.local
- Default Gateway:** 10.48.40.1
- DNS Server 1:** 8.8.8.8
- DNS Server 2:** 134.141.121.33

Figure 8: Network Setup - Host Attributes

3. Select additional ports as necessary to display and verify the interface settings. If you make any changes to the additional interface settings, select **Save**.



Note

If you want to configure a port interface, or add a VRRP alias that can be used by an engine application, go to [Add a Port](#) on page 29.



Note

Configuring an engine instance IP address for the admin interface of the container application is useful for diagnostic purposes. The web interface can be accessed from the [Engine Settings](#).

Add a Port

Use this optional procedure to configure the port interface on one of the Universal Compute Platform data ports. You can use this configuration to set up a VRRP IP alias on the port.




Note

Some container applications require you to configure a VRRP IP alias on one of the Universal Compute Platform data ports. The same VRRP address can be assigned to the data port on Universal Compute Platform, and to the engine application to provide an IP address that lets you access the application's management interface using a web browser.

To add a new port interface, take the following steps:

1. Navigate to **Administration > System > Network Setup**.
2. From the **Interfaces** pane, select **Add New Interface**.
3. Enter a **Name** and **IP Address** for the port interface.
4. Configure the additional fields in the Port settings window. For help with the fields and their settings, see [Create New Interface Settings](#) on page 30.
5. Optional. To create an interface alias using VRRP, do the following:
 - a. Under **VRRP**, select **Add**.
 - b. Enter the VRRP **Address** along with optional **Comments**.
 - c. To add additional addresses, repeat the previous two substeps.

You can add up to 10 VRRP addresses on a single data port. To delete an address, select the adjacent  icon.
 - d. Enter a **Priority** and **Router ID** that applies to the VRRP configuration on this port.
6. Select **Save**.



Note

Later on, when you install the application engine, complete the IP alias configuration by assigning the same VRRP address to the **Assigned Virtual IP Address** field for the application.


Create New Interface Settings

The Universal Compute Platform enables the creation of additional logical interfaces to support multi-homed deployments. Use the **Create New Interface** wizard to configure the required settings. The following table describes the settings.

Table 10: Interface Properties

Field	Description
Name	Name of the interface.
Mode	Describes how traffic is forwarded on the interface topology. Options are: <ul style="list-style-type: none"> 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.
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.
Layer 3	
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 10: Interface Properties (continued)

Field	Description
FQDN	Fully-Qualified Domain Name
VRRP	<p>You can assign VRRP settings to create an IP alias that provides access to the underlying management interface for an application instance. To add a VRRP address, select Add and configure the following:</p> <p>Address</p> <p>Enter the VRRP IP address that you want to assign to the management interface of the application instance.</p> <ul style="list-style-type: none"> • VRRP IPs must be in the same network segment as the data interface IP. • VRRP IPs must not overlap with other addresses. • You can add up to ten VRRP addresses to a data interface. To delete an address, select the adjacent  icon. <p>Comments</p> <p>An optional text comment that identifies the purpose of the VRRP address. You can enter a separate comment for each VRRP address.</p> <p>Priority</p> <p>A numeric value (between 1 - 254) that is used to arbitrate mastery of the state of exchanges for members of a cluster. The priority that you enter applies to all VRRP addresses on this port.</p> <p>Router ID</p> <p>A numeric value (between 1- 255) that is used to segment a routing domain. The router ID that you enter applies to all VRRP addresses on this port.</p> <p>Note: In a Self-Orchestration standalone cluster, the specific values that you enter for Priority and Router ID are not important. These settings are important in multi-node configurations only.</p>

Install Kubernetes (Offline Installation)

Use this procedure only if you are installing and configuring the standalone cluster while offline. In this case, you must upload and install a Kubernetes image before you create the cluster.




Note

For online installations, you can skip this procedure as Kubernetes installs automatically when you create the standalone cluster.

1. Go to **Administration > System > Software Upgrade**.
2. Select the **Kubernetes** tab.
3. Upload the image file using one of the following two methods:
 - On the **Kubernetes** page, select **↑ (Upload Kubernetes resource package zip file)** and then browse to the image file and select it.
 - On your desktop, select the image file and drag it into the **Upload Kubernetes resource package zip file** area of the **Kubernetes** page.

The image file uploads to Universal Compute Platform and then installs automatically.

**Note**

If you want to delete an Available image, select the adjacent  (Remove Mirror).

Configure the Stand-Alone Cluster Settings

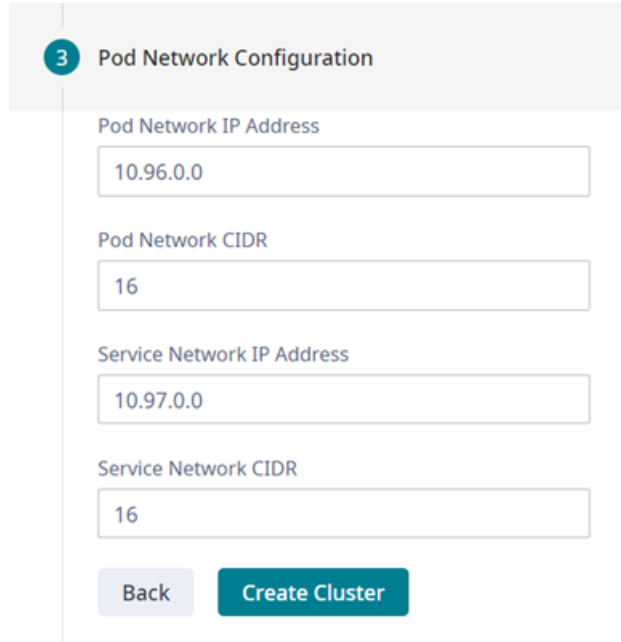
An engine is an instance of a containerized application. This process follows the user interface to configure the orchestration engine settings for a stand-alone deployment. From the management IP address, log into the user interface using the admin credentials that you configured using the Basic Configuration Wizard.

Go to **Cluster Settings > Cluster Configuration** and configure the stand-alone cluster in the following order, as shown on screen:

1. Deployment Type
2. Cluster Mode
3. Pod Network Configuration
4. Finish

To configure the cluster, do the following:

1. For the **Deployment Type**, select **ExtremeCloud Edge - Self-Orchestration**.
2. In the **Cluster Mode** section, select **Standalone** and click **Next**.
3. Provide the settings for **Pod Network Configuration**:
 - 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)



3 Pod Network Configuration

Pod Network IP Address
10.96.0.0

Pod Network CIDR
16

Service Network IP Address
10.97.0.0

Service Network CIDR
16

Back Create Cluster

Figure 9: Pod Network Default Settings

4. Select **Create Cluster**.
5. Select **Done**.




Note

The cluster state is bound to the IP Address of the ICC interface. Do not change the ICC IP address once the cluster is set. If the ICC IP address changes after the cluster is created, the cluster state resets (even for standalone clusters) resets. If this occurs, cluster configuration must be re-initialized and any installed applications must be re-installed.



Note

For online installations, Kubernetes installs automatically when you create the cluster.



Engine Application Installation

[Download the Docker Application Image](#) on page 35

[Upload Application Image to Appliance](#) on page 35

[Install Engine Application](#) on page 35

[Deploy Application](#) on page 35

[Configure Interface Settings for Engine Application](#) on page 36

[Engine Upgrades](#) on page 37

To install an engine application, complete the following tasks.



Note

Refer to your engine application documentation for requirements that are specific to the engine.

Table 11: Engine Application Installation

Step	Procedure	Description
1	Download the Docker Application Image on page 35	Download the application image from the Extreme Networks support portal.
2	Upload Application Image to Appliance on page 35	Upload the application image to the Universal Compute Platform appliance.
3	Install Engine Application on page 35	Install the application image on the appliance.
4	Deploy Application on page 35	Deploy the application engine on the appliance.
5	Configure Interface Settings for Engine Application on page 36	Configure application interface settings specific to the application. For example, you can obtain the Locking ID for licensing or assign a VRRP IP alias.
6	Engine Upgrades on page 37	Upgrade the application software. Select the upgrade method that fits your application type.

What to do Next

To onboard the application to the cloud, go to [Onboard Cluster to ExtremeCloud IQ](#) on page 39.

Download the Docker Application Image

To obtain the application 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 `rtxcx-9000.tar`.

Upload Application Image to Appliance



Upload the Docker application image file to the Universal Compute Platform appliance.

1. Go to **Engines > Image Management**.
2. Complete either of the following options:
 - Select the **Choose Image File** pane and navigate to the image file.
 - Drag and drop the image file onto the **Image File** pane.

A list of uploaded image files displays below the **Choose Image File** pane.



Note

To delete an uploaded image, select the check box next to the image file. Then, select . To refresh the image file list, select .

Install Engine Application

Install the Docker application image file on the appliance.

1. Go to **Engines > Installation**.
2. From the application pane for the intended application, select **Install**.



Note

- If you have not yet uploaded the application Docker image file, you will be prompted to do so.
- The installation time will depend 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

After you have installed the engine on the appliance, deploy the application.

1. Go to **Engines > Installation**.
2. Select the engine instance. link. For example, "cx9000 #1".
3. Assign virtual functions (VFs) to the ports. We recommend that you retain the default settings, although you can assign any unassigned VF.
4. Select **Deploy**.

Configure Interface Settings for Engine Application

Use this procedure to configure the final application interface settings immediately after 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 36.

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 screen (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 Platform 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**.

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 38.
- **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.



Onboard Cluster to ExtremeCloud IQ

[Onboarding a Cluster to ExtremeCloud IQ](#) on page 39

[Cloud Visibility](#) on page 40



Note



For Self-Orchestrated deployments, onboarding to ExtremeCloud IQ is optional. Use the topics in this section only if you plan to onboard to ExtremeCloud IQ.

After the Universal Compute Platform cluster is installed, associate the node cluster with your ExtremeCloud IQ account:

1. Onboard the cluster to your ExtremeCloud IQ account. See [Onboarding a Cluster to ExtremeCloud IQ](#) on page 39.
2. Onboard your devices and operate the account.

Onboarding a Cluster to ExtremeCloud IQ

To onboard a Universal Compute Platform cluster into ExtremeCloud IQ use the ExtremeCloud IQ Quick Add function:

1. From the ExtremeCloud IQ main navigation pane, select  (Manage), and then select **Devices**.
2. Select  (**Add**) and then select **Quick Add Devices > Manage your devices directly from the cloud**.

- In the **Serial Number** field, enter the serial number for one node in the cluster.

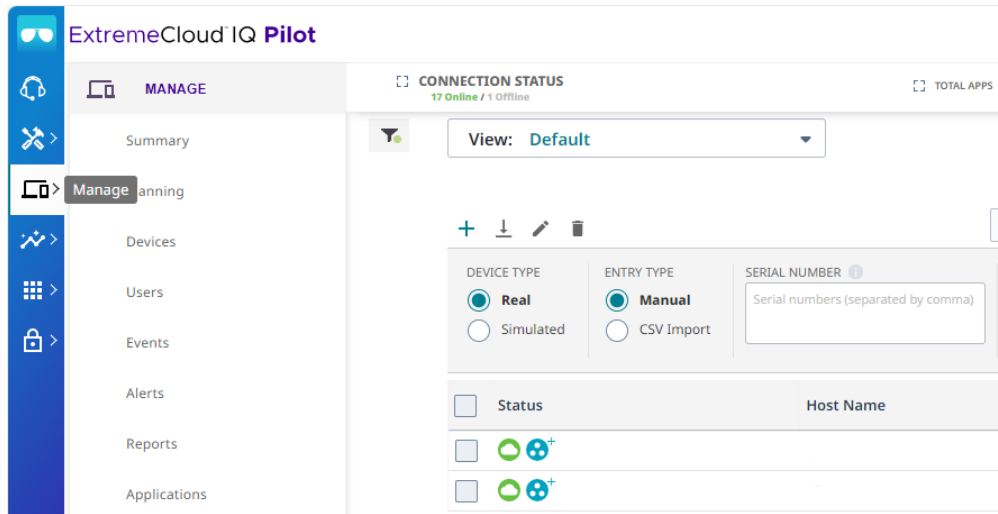


Figure 10: Add Cluster to ExtremeCloud IQ

The **Device Make** field displays.

- From the **Device Make** menu, select **Universal Appliance**.
- Select **Add Devices**.

The full cluster is added based on the serial number of a single node in the cluster.



Note


To view details about the cluster, select the **Host Name** link.

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.



Index

A

- admin password configuration 21
- announcements 8, 9
- application support 12

B

- basic configuration wizard
 - run the 21

C

- cloud visibility
 - get cloud address 40
- cluster
 - configure a standalone cluster 32
 - onboard to cloud overview 39
 - onboard to ExtremeCloud IQ 39
 - prerequisites 10
- configuration
 - assign reserved IP ranges 32
 - configure node IP address 32
 - configure self-orchestration 20
 - select deployment type 32
 - validate the network address configuration 27
- console port
 - connect to appliance using 21
- conventions
 - notice icons 6
 - text 6

D

- data port settings
 - assign port and vlan 21
 - management on data interface 21
- deploying Stand-alone Universal Containers 10
- deployment overview 10
- documentation
 - feedback 9
 - location 7, 8

E

- engine application
 - assign VRRP alias to 36
- engines
 - deploy application 35
 - download docker application image 35

- engines (*continued*)
 - engine application settings 36
 - install application image on appliance 35
 - installation workflow 34
 - upgrade application 37
 - upload application image to appliance 35
- example of internal networking 14
- ExtremeCloud IQ
 - firewall requirement for connection to 18
 - get cloud address 40
 - onboard cluster to cloud 39
 - onboard cluster to cloud overview 39

F

- feedback 9
- firewall requirements
 - for connection to ExtremeCloud IQ 18
 - for online installs and upgrades 17

G

- global default gateway settings 21

H

- hardware support and specifications 12
- high availability 16
- host attribute settings
 - set hostname, domain name, and DNS 21

I

- ICC
 - connect to appliance using ICC1 port 21
 - ICC port settings 21
 - IP address required but not connectivity 10, 21
 - kubernetes binds to 13
- install kubernetes (offline) 31
- installs
 - firewall requirements for online 17
- interfaces
 - create a new 30
 - interface properties 30
- internal networking example 14
- internet connectivity requirements 17
- IP addresses
 - reserved ranges 18

K

kubernetes
binding to ICC address 13
offline installation 31
online installation 32
overview for self-orchestration 13
reserved IP ranges for 18

N

network architecture 14
notices 6

O

OS access 14

P

port requirements
requirements for online installs and upgrades 17
ports
adding a 29
product announcements 8, 9

R

requirements
reserved IP ranges 18
when deploying VRRP 18

S

support
hardware and application support 12
technical support 8, 9

T

technical support
contacting 8, 9
test connectivity using ping 21
time settings
set NTP servers and time zone 21

U

upgrades
engine upgrade overview 37
firewall requirements for online 17
upgrade self-orchestrated application 38
upgrade the appliance 25

V

virtual switch overview 14
vrrp
assign vrrp alias to application 36
VRRP

VRRP (*continued*)

assign VRRP alias to application 36
VRRP (Virtual Router Redundancy Protocol)
alias requirement for some applications 29
configuration 18
not required on ICC 21

W

warnings 6