Extreme networks

ExtremeCloud Edge v5.13.01 Managed Orchestration Deployment Guide

Configuration and Management from Universal Compute Platform

9039472-00 Rev. AA
December 2025

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/

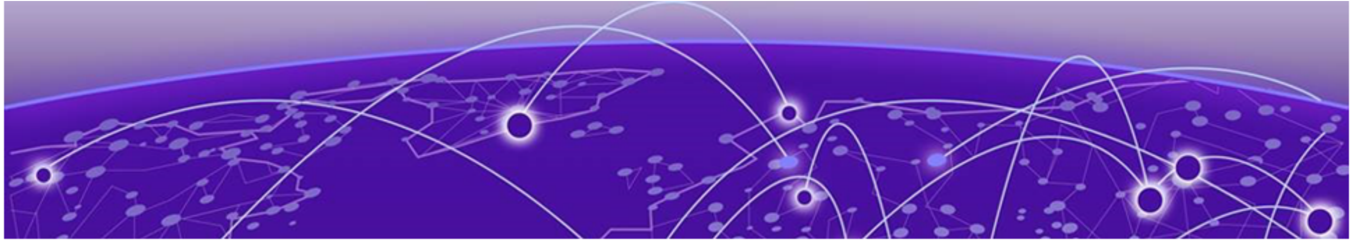

Table of Contents

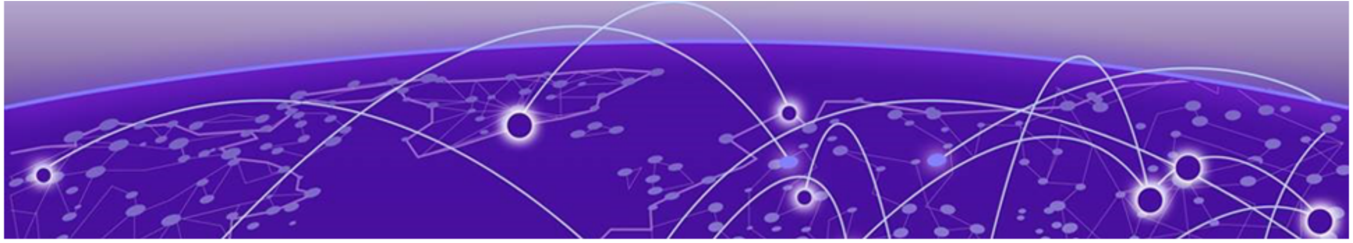

Abstract

The ExtremeCloud Edge v5.13.01 Managed Orchestration Deployment Guide provides comprehensive procedures for deploying ExtremeCloud Edge within the Universal Compute Platform environment. Key setup requirements for Managed Orchestration include detailed hardware specifications and network configurations, such as Virtual Router Redundancy Protocol (VRRP) for high availability and load distribution. The guide emphasizes the deployment of Kubernetes clusters and Docker containers for application orchestration, along with configuring firewall policies and network addressing schemes. It delves into the intricacies of inter-cluster communication for synchronization and node health monitoring over a 10 Gbps backplane. Additionally, the guide covers the configuration of cluster states, pre-deployment readiness checks, and software provisioning, including the integration with the ExtremeCloud IQ engine for centralized cloud management. Detailed instructions on network segmentation, private IP schemes, and persistent connections to ExtremeCloud services ensure robust operational performance.

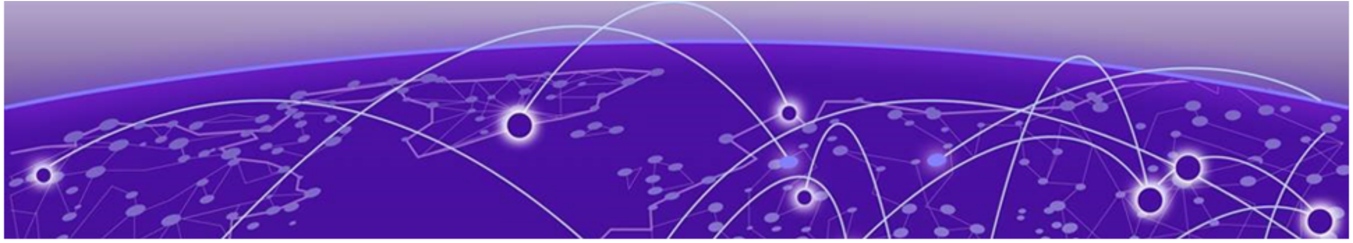

Preface

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.

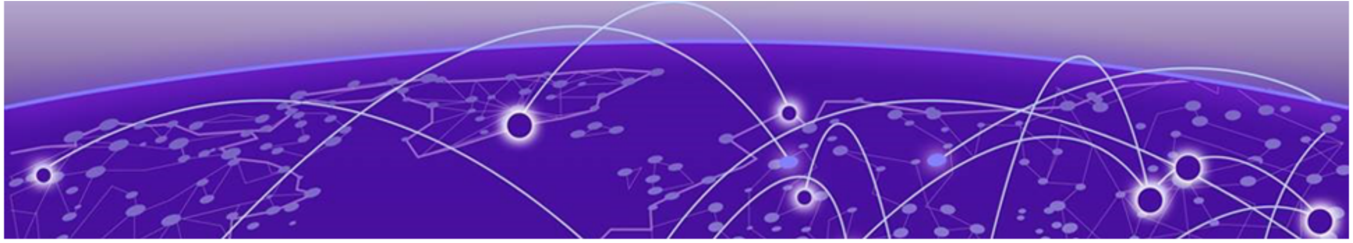

Introduction

The *ExtremeCloud Edge - Managed Orchestration Deployment Guide* provides comprehensive procedures for deploying and configuring the infrastructure that enables the running of ExtremeCloud applications at a customer's premises, as supported by the ExtremeCloud Edge - Managed Orchestration offering.

This document details the steps related to setting up the hosting environment, which consists of a variable-sized cluster of Universal Compute Platform hosts and the corresponding network and access requirements for the solution. The Universal Compute Platform cluster provides a Kubernetes-based containerized environment over which the components (container microservices) of the ExtremeCloud application portfolio are installed. The cluster essentially abstracts the local installation to mirror the platform and middleware dependencies of ExtremeCloud public installation.

This guide outlines the installation steps and setup for the hosting environment, in advance of the installation of ExtremeCloud application suites. Installation and management of the lifecycle of ExtremeCloud application software is provided by ExtremeCloud OPS as a managed service.

Network setup and Access configuration are key elements of the installation. This document provides configuration details pertaining to the setup of the cluster and the application interface to the network. The setup includes network addressing, configuring the Virtual Router Redundancy Protocol (VRRP), and crucial firewall settings. VRRP support provides high availability and load balancing while the firewall settings are required for egress and ingress of service operations as well as CloudOPs lifecycle management.

Additionally, the guide covers the configuration of cluster states, pre-deployment readiness checks, and software provisioning. This includes integration with the ExtremeCloud IQ engine for centralized cloud management. This guide provides detailed information about network segementation, private IP schemas, and persistent connections to ExtremeCloud IQ services.

Related Links

Managed Orchestration Deployment Training Video

Managed Orchestration Cluster

Universal Compute Platform leverages Kubernetes and Docker to deploy and manage the delivery of applications to the customer premises.

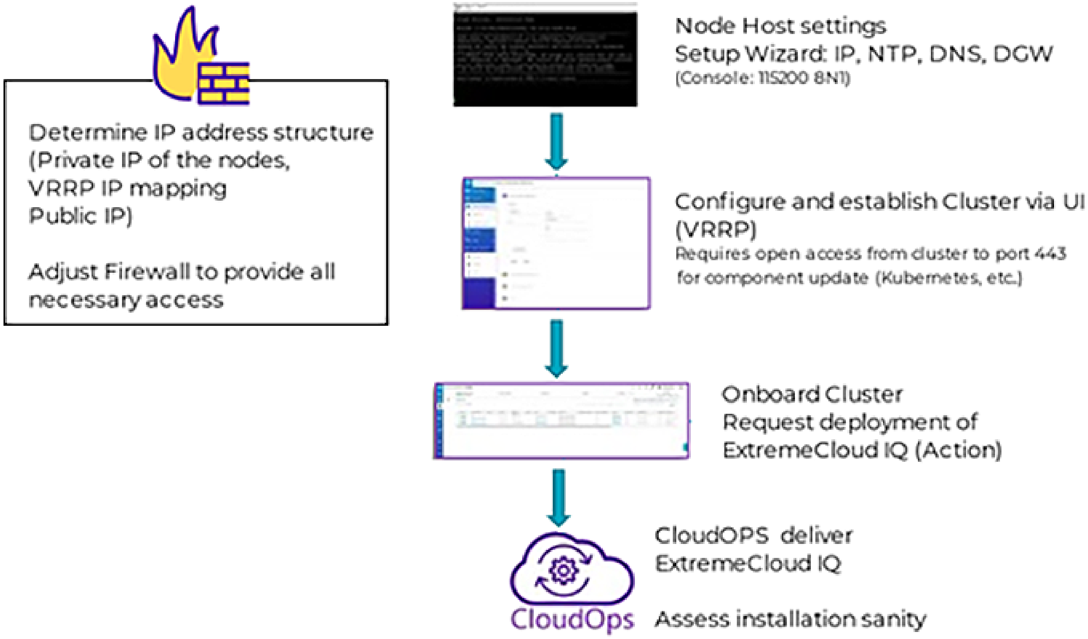
Determine IP address structure
(Private IP of the nodes,
VRRP IP mapping
Public IP)
Adjust Firewall to provide all necessary access
Node Host settings
Setup Wizard: IP, NTP, DNS, DGW
(Console: 115200 8N1)
Configure and establish Cluster via UI (VRRP)
Requires open access from cluster to port 443 for component update (Kubernetes, etc.)
Onboard Cluster
Request deployment of ExtremeCloud IQ (Action)
CloudOPS deliver ExtremeCloud IQ
Assess installation sanity
CloudOps

Figure 1: ExtremeCloud IQ Deployment Workflow

The following figure depicts the three physical host boxes required for Universal Compute Platform, with ports mapped as follows:

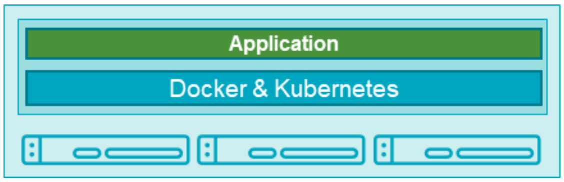

As an option, the system leverages VRRP (Virtual Router Redundancy Protocol) to provide support for both high-availability and load balancing, supported by an NGINX engine. All service operations to the cluster should be directed to the corresponding VRRP IP so that the load balancing logic can direct the request to the best node.

Deployment configuration requirements vary over different applications deployed into the Universal Compute Platform. One main requirement in the establishment and operation of the cluster is the Inter-Cluster Connection. This connection operates as the backplane between nodes in the cluster. This backplane carries all the synchronization data between nodes for both component and data states. It is a best practice to deploy the interface as a segregated 10 Gbps inter-connect (separate switch port), allowing for the best performance in synchronization between nodes.

- Inter-Cluster Connection: Backend interaction and synchronization between all the members of a cluster. Minimum required connection requires 10 Gbps between nodes.
- The internal Kubernetes engine requires the reservation of two (2x) /16 subnets. This set of IP Address ranges is for internal use only by inter-component and framework operations. This reserved range can be anything, but customer should ensure that this IP address range does not conflict with any routable address space within the organization.

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 components must be downloaded and installed during the cluster configuration stage. After you select **ExtremeCloud Edge – Managed Orchestration** as the cluster type and initialize the cluster, the appliance connects to Docker Hub to download additional Kubernetes components based on your installation requirements. The appliance installs the components and creates the cluster.

For example, the installation may reach out to Docker Hub and redirect to `registry.k8s-io` and `amazon.aws` as follows: `https://prod-registry-k8s-io-eu-central-1.s3.dualstack.eu-central-1.amazonaws.com`.

After the cluster is created, Kubernetes binds to the ICC VRRP address.

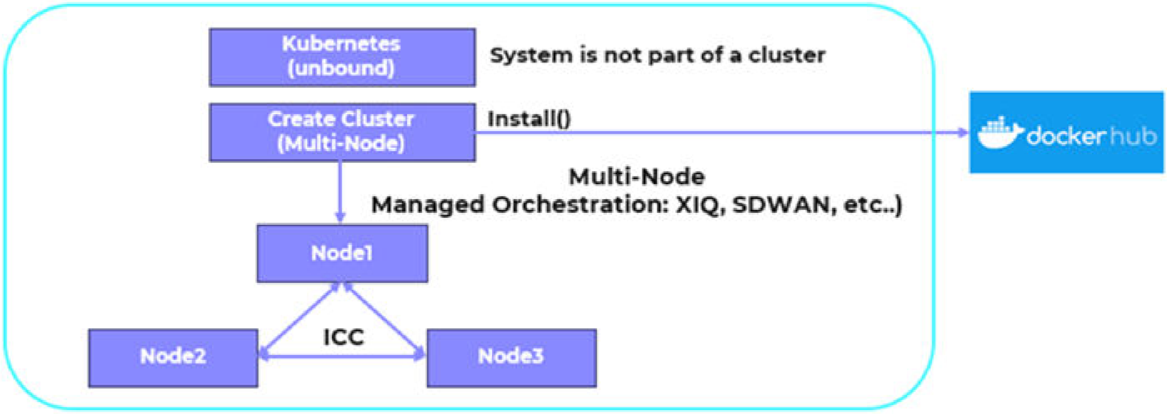
Kubernetes (unbound)
System is not part of a cluster
Create Cluster (Multi-Node)
Install()
docker hub
Multi-Node
Managed Orchestration: XIQ, SDWAN, etc..)
Node1
ICC
Node2
Node3

Figure 2: Cluster Creation with Kubernetes

Note

- Because of the ICC binding, it's recommended to use the data ports for application management rather than the ICC ports. If you change the ICC address or ICC VRRP address, the Kubernetes binding breaks and the Kubernetes installation unwinds, effectively wiping out the installation with the fix being to reinstall and reconfigure.
- 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.

Deployment Responsibilities

This topic outlines the key deployment responsibilities for deploying ExtremeCloud Edge on Universal Compute Platform. Each of the following key personnel have unique responsibilities:

- Customer On-Site Representative
- Extreme CloudOps
- System Administrator of Universal Compute Platform

Customer On-Site Representative

Customer On-Site Representatives are responsible for the following tasks:

- Set up a firewall that enables cluster access to the appropriate internet ports (for example, port 443) and enables CloudOps access. Follow the firewall configuration guidelines under Firewall Access for Critical Settings on page 18.
- Configure each node for service — Provide the necessary IP, DNS, and Host addresses, ICC Configure and form cluster (VRRP).
- Register the cluster with an ExtremeCloud IQ Public account.
- Register an ExtremeCloud IQ deployment request. The request requires a valid XIQ-EDGEOPS-S-EW in good standing. This SKU is a required component of an ExtremeCloud Edge BOM quote.

For detailed information, see Managed Orchestration Cluster on page 11.

Note
The *ExtremeCloud Edge - Managed Orchestration Deployment Guide* covers tasks that are mostly completed by the Customer On-Site representative.

Extreme CloudOps

ExtremeCloud IQ CloudOps is responsible for the following tasks:

- Deploy ExtremeCloud applications to the Universal Compute Platform cluster.
- Create monitoring and backup frameworks.
- Validate the state of all operational components.

Universal Compute Platform Administrator

Universal Compute Platform Administrators are responsible for the following tasks:

- Create ExtremeCloud IQ user accounts for end-device management.
- Onboard managed devices from the ExtremeCloud IQ local account.

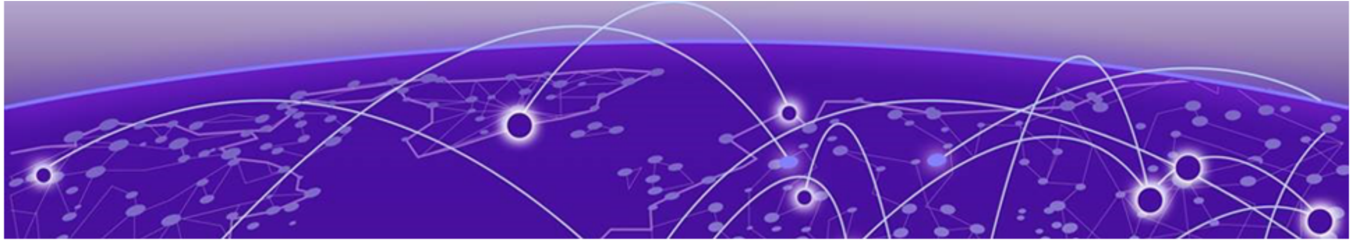

Requirements

This section outlines requirements for deploying a Managed Orchestration Deployment of ExtremeCloud Edge.

ExtremeCloud Edge Planning

This guide outlines the steps required to prepare a cluster environment that will support deployment of ExtremeCloud Edge applications to the customers' premises.

Minimum Requirements for Installation

- Five Public IP addresses exposed via the firewall and port-forwarded to the internal service sets
- Firewall adjustments to allow communication of system functions to external entities (licensing, component upgrades, device management) and CloudOPS access for lifecycle management of the intalled applications/software. Please refer to section Firewall Access for Critical Settings on page 18.
- A cluster of Universal Compute Platform appliances.

  Note
  The cluster size must be a multiple of three. The minimum cluster size for ExtremeCloud IQ with up to 5,000 devices is three nodes. However, six nodes is the typical size for most deployments. Check with your sales representative to size your deployment according to your application choices and capacity requirements.

- Network Connectivity for the hosts both in backplane (ICC) and application data operations (data ports). 10 Gbps minimimun links recommended.
- ICC: Interconnect (backplane) for cluster operations, component state and shared filesystem synchronization. Each node requires connection of ICC to common backplane network segment.
- Data: Interfaces that the applications will use with other devices or systems for operation management, such as remote device management (for access points

and switches) and license services. Data interface is also utilized for remote lifecycle management of installed software.

Application requirements for the cluster configuration:

- Five IP addresses representing the various services offered by the application to provide load balancing (Service Set 1 – 4).
- Each node in the cluster must map each of the services to a data interface, and all services can be mapped into the same interface. The same data interface can represent a direct point of reference for each of the front-end VRRP services.
- Five VRRP IP address are required to support port-overlap services for different services or a functional model (such as CAPWAP Master vs CAPWAP Server).

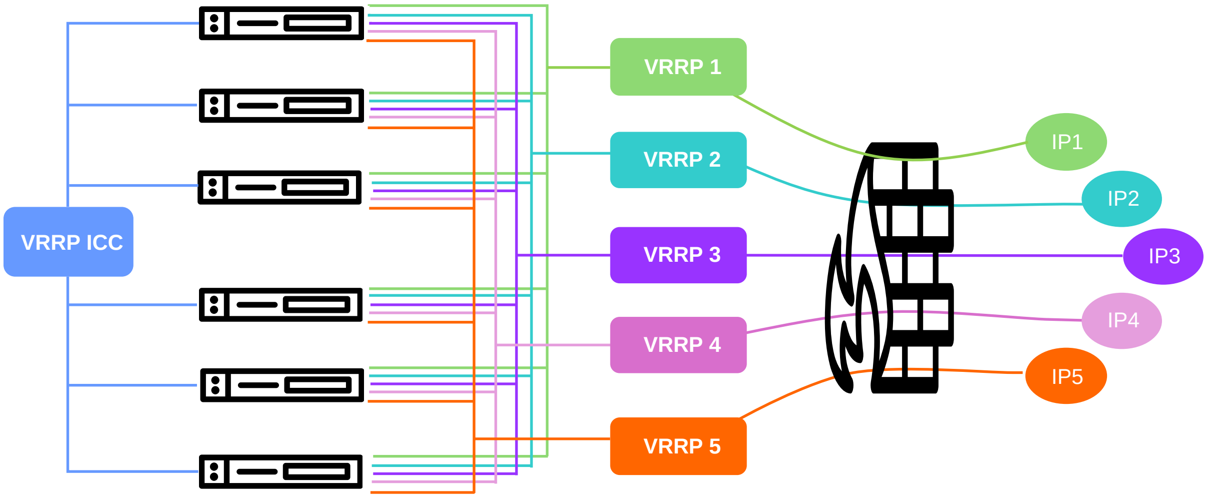

Out-of-Bound Routing for Outgoing Traffic

The VRRP service set mappings provide load balancing and service abstraction for incoming traffic. For outgoing traffic that originates from installed components, including responding to incoming traffic that came through these mappings, is steered through the default gateway. At the internal firewall, the source address for the outgoing traffic is the address of the data interface on the node from which the traffic originated.

Service Set 1: Administration, Account Access (https), CAPWAP Master, Diagnostics

Table 4: Example port assignments for Service Set 1

Port	Protocol	Service	Description
80	TCP	CAPWAP	CAPWAP Master
443	TCP	NGINX	ExtremeCloud IQ Admin, software management
1443	TCP	XAPI	ExtremeCloud IQ API
2083	TCP	IDM	IDM Auth
12222	UDP	CAPWAP	CAPWAP Master

Service Set 2: AP Registration/CAPWAP Load Balancing

Table 5: Example port assignments for Service Set 2

Port	Protocol	Service	Description
80	TCP	CAPWAP	CAPWAP Master
443	TCP	SD-WAN	SD-WAN Communicator
5825	TCP	Inlets	Device Communication
8090	TCP	Inlets	Device Communication
9090	TCP	SD-WAN	SD-WAN Communicator
12222	UDP	CAPWAP	CAPWAP Master

Service Set 3: AP Registration/CAPWAP Load Balancing

Table 6: Example port assignments for Service Set 3

Port	Protocol	Service	Description
80	TCP	CAPWAP	CAPWAP Master
443	TCP	SD-WAN	SD-WAN Communicator
12222	UDP	CAPWAP	CAPWAP Master

Service Set 4: AP Registration/CAPWAP Load Balancing

Table 7: Example port assignments for Service Set 4

Port	Protocol	Service	Description
80	TCP	CAPWAP	CAPWAP Master
443	TCP	NGINX	ExtremeCloud GDC access
1443	TCP	XAPI	ExtremeCloud IQ API
1444	TCP	API	GDC API
1445	TCP	License	ExtremeCloud IQ License Management
12222	UDP	CAPWAP	CAPWAP Master

Service Set 5: Login Redirection

Table 8: Example port assignments for Service Set 5

Port	Protocol	Service	Description
443	HTTPS	Login	Extreme Platform One Login Rediection

Firewall Access for Critical Settings

In a typical on-premise installation, the cluster is installed behind an access firewall, providing network address translations between the public and private address spaces. Always allow access for CloudOps management of the cluster. The standard deployment of ExtremeCloud Edge requires five public IP addresses to front-end the installation. They are mapped to forward traffic into the five VRRP IP addresses of the service sets.

During system setup, the following configuration settings are critical to the deployment:

- Default Gateway: Each node in the cluster supports a single default gateway (0.0.0.0/0) definition. This gateway must be mapped to a next-hop attached on the data port interface.

  Note
  Do not configure the default gateway to map to the Inter-Cluster Connection (ICC) interface. The ICC is an internal connection between systems that is not used for management or operation of the cluster.

- DNS server: At least one reachable DNS server must be configurable, allowing the system to resolve several URLs during installation and interaction with ExtremeCloud IQ and CloudOps functions.
- Network Time Protocol (NTP) Servers: At least one reachable NTP, allowing the system to synchronize its time with a trusted time source. The same NTP must be configured, in the same order, on all nodes in the cluster.

  A best practice is to have two NTP definitions to support availability of the primary server. If there is an issue with the primary server, the system resorts to the alternate server.

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.

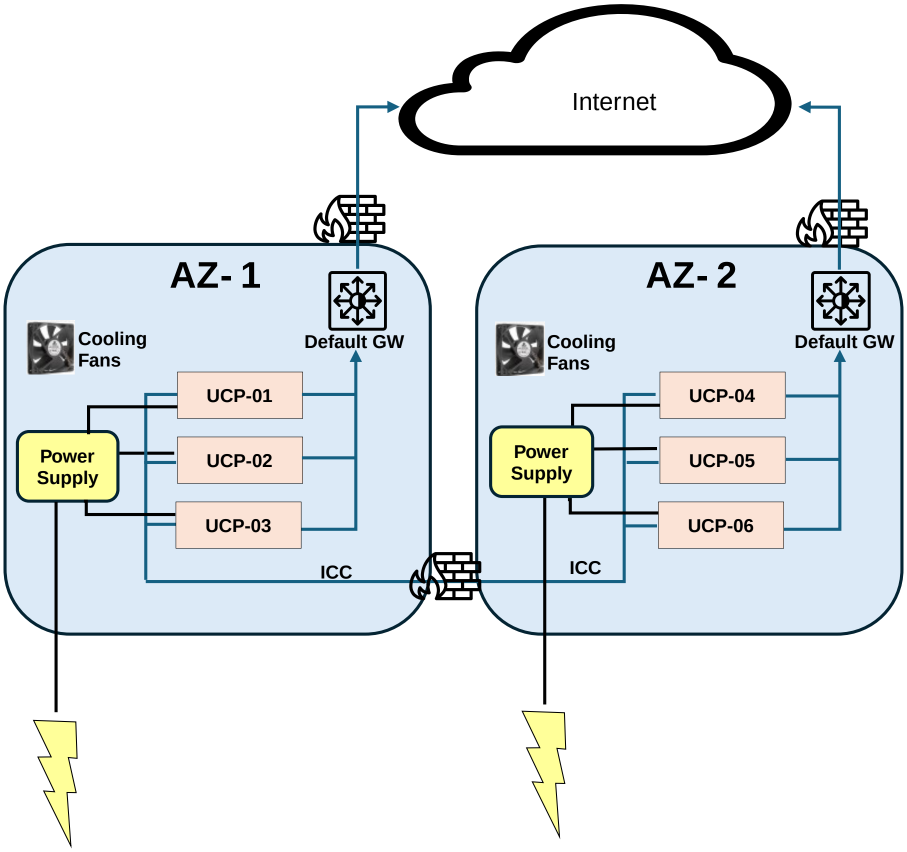

Figure 3: Multiple Availability Zones for a Six-Node Cluster

Prerequisites for ExtremeCloud Edge Installation

Address planning is the fundamental step in successful deployment of the Universal Compute Platform to support installing ExtremeCloud applications such as ExtremeCloud IQ. It is important to understand the following:

- Decide how you will deploy and access the services offered by the cluster. Is the cluster going to serve applications that operate only within the on-premises installation? Or is application access going to require external access? Pre-determination of the IP address and connectivity structure are fundamental to a successful deployment. These deployment decisions drive the configuration choices.
- Consider the address plan of the installation, including how the cluster is going to be presented externally via a firewall.

  Each externally exposed address must be mapped to an internal VRRP of the cluster. You can either directly expose the VRRP IP addresses for the five service sets directly through a firewall, or in the case of NAT translation, ensure that the externally available IP addresses are mapped 1:1 to the internal services, and that the correponding application ports are allowed for access (per firewall rules definition).

Important

Before you begin step-by-step configuration, make sure that you clearly understand and document all the elements of the network presence and topology related to the deployment.

The Inter-Cluster Connection (ICC) IP address is critical to the continuous operation of the system. If address definitions for ICC require re-addressing, the entire cluster will need to be rebuilt and the application re-deployed in order to re-establish all the correct references of services within the cluster.

It is strongly recommended that the *entire* IP address structure for all services be defined once and not changed. Re-addressing can expose internal dependencies on references to mapped services and therefore affect the integrity and stability of the deployed installation.

IP Addresses

The most important point of definition is to record the IP address relationship between the cluster’s direct interfaces (Node, Service Set, Virtual IP address (VIP)), and external access. Each node has it's own data interface IP address.

Table 9: IP address relationship between the cluster's direct interfaces and external access

Service Set	Virtual IP (VIP)	Public IP
Service Set 1 (cmudp, cmtcp, cmauth, https)	*VIP 1*	*Public IP 1*
Service Set 2 (csupd1, cstcp1)	*VIP 2*	*Public IP 2*
Service Set 3 (csudp2, cstcp2)	*VIP 3*	*Public IP 3*

Table 9: IP address relationship between the cluster's direct interfaces and external access (continued)

Service Set	Virtual IP (VIP)	Public IP
Service Set 4 (csudp3, cstcp3)	*VIP 4*	*Public IP 4*
Service Set 5 (https)	*VIP 5*	*Public IP 5*

VRRP Configuration

In support of load balancing and high-availability functions, the Universal Compute Platform relies on Virtual Router Redundancy Protocol (VRRP) to provide IP abstraction to key functionality. VRRP is critical in the configuration model.

The following operation settings must be defined as part of the VRRP configuration of member nodes:

- **Priority**— VRRP uses priority settings as a mechanism to arbitrate mastery of the state of exchanges across members of the cluster.

  The node with the higher priority defaults to the master. However, in the case of failovers of the master node, VRRP algorithms assign mastery to the next higher priority member of the cluster. Therefore, it is important to properly assign corresponding priority settings to each node, so that their hierarchical priority in terms of VRRP state ownership is clear.

  As a best practice:

  - Designate node 1 as the highest priority, node 2 for second highest priority, and nodes 3-6 as lower priority.
  - The same priority should be used across all services (ICC, Services)
- **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.

Inter-Cluster VRRP Configuration

An Inter-Cluster Connection refers to the back-end interaction and synchronization between all the members of a cluster. Minimum required connection requires 10 Gbps between nodes.

Table 10: Inter-Cluster Connection VRRP Configuration

	Nodes 1 -6 (Port #)
ICC	• Node 1 ICC IP /CIDR • Node 2 ICC IP/CIDR • Node 3 ICC IP/CIDR • Node 4 ICC IP/CIDR • Node 5 ICC IP/CIDR • Node 6 ICC IP/CIDR
VLAN	VLAN Tagged/Untagged
Port type	Physical
VRRP	
VRRP IP addresses	ICC VRRP IP
Priority	Set a unique priority for each node. For example: • Highest (200) • Next (150) • Medium (100) • Next (75) • Next (50) • Low (25)
Router ID	ID (2)

Services VRRP Configuration

The VRRP configuration relates to the number of services you are exposing. Configure a VRRP IP address (VIP) for each service.

Table 11: Services VRRP Configuration

	Nodes 1-6 (Port #)
Data Port	Node Port IP /CIDR. Unique Port IP for each node.
VLAN	VLAN Tagged/Untagged
Port type	Physical
VRRP	
VRRP IP address (VIP)	6 VIP addresses. Unique VIP for each node
Priority	Unique priority value for each VIP
Router ID	ID (1)

Reserved IP Addressing

Non-routable container orchestration by Kubernetes within the cluster requires reservation of private non-routable 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. Make sure to use private non-routable address spaces for any ranges that you adjust.

Table 12: 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.

Port Information for Firewalls

Map the following service ports to the Service Set VRRP IP addresses listed in IP Addresses on page 20.

- VLAN/VIP address for CAPWAP Master and API services (TCP 80/UDP 12222/TCP 2083/443)
- VLAN/VIP address for CAPWAP Server 1 service (TCP 80/UDP 12222)
- VLAN/VIP address for CAPWAP Server 2 (TCP 80/UDP 12222)

ExtremeCloud IQ on-premises installations require access to ExtremeCloud IQ core services. Make sure the firewall configuration allows for access to ExtremeCloud IQ core services.

The following tables list outbound ports for use when the firewall configuration requires rules that enable outbound traffic.

Basic Access for ExtremeCloud Services

This is required for ExtremeCloud applications to run properly on ExtremeCloud Edge RDC.

Note

The Readiness Assessment tool requires access to ucp0-console-extremecloudiq.com at 18.192.120.159 during initial deployment. Run the Readiness Assessment tool before you install the ExtremeCloud IQ engine. The tool accesses this server to test your firewall configuration and to assess whether your system is ready for the ExtremeCloud IQ engine installation.

Table 13: Firewall Configuration Details (Outbound Traffic)

Domain Name	IPv4 Addresses	Protocol	Port
hac.extremecloudiq.com	34.253.190.192 ~ 34.253.190.255	HTTPS	443
<rdc>-inlets.extremecloudiq.com	Dynamic IP range	TCP	8090
hmupdates-ng.aerohive.com	54.86.95.132	HTTPS	443
extremecloudiq.com	34.253.190.192 ~ 34.253.190.255	HTTPS	443
	18.194.95.0 ~ 18.194.95.15		
	3.234.248.0 ~ 3.234.248.31		
	44.234.22.92 ~ 44.234.22.95		
prod-mwapi.extremenetworks.com	134.141.117.210	HTTPS	443
informaticacloud.com	Dynamic	HTTPS	443
mx.extremecloudiq.com	34.202.197.56/57	TCP	587
stun.extremecloudiq.com	3.234.248.28 - 29	UDP	12222
api.ip2location.com	Dynamic IP range	HTTPS	443
docker.io	Dynamic IP range	HTTPS	443
gcr.io	Dynamic IP range	HTTPS	443
k8s.io	Dynamic IP range	HTTPS	443
maven.org	Dynamic IP range	HTTPS	443
Amazon S3	Dynamic IP range	HTTPS	443
amazonaws.com	Dynamic IP range	HTTPS	443
NTP Service	<Any NTP Server IP>	UDP/TCP	123
extremeportal.force.com	Dynamic IP range	HTTPS	443
prod.extreme.sentinelcloud.com	Dynamic IP range	HTTPS	443
cloud-status.extremecloudiq.com	18.67.39.6	HTTPS	443
cloud-cdn2.extremecloudiq.com	Dynamic IP range	HTTPS	443
rest.nexmo.com	Dynamic IP range	HTTPS	443
extremesaas.my.site.com	Dynamic IP range	HTTPS	443
charts.xcloudiq.com	Dynamic IP range	HTTPS	443

Table 13: Firewall Configuration Details (Outbound Traffic) (continued)

Domain Name	IPv4 Addresses	Protocol	Port
github.com	Dynamic IP range	HTTPS, SSH	443, 22
vault-prod.extremecloudiq.com	Dynamic IP range	HTTPS	443
cache.extremecloudiq.com	Dynamic IP range	HTTPS	6379
events.pagerduty.com	Dynamic IP range	HTTPS	443

Access

Table 14: Outbound Traffic

Domain Name	IPv4 Addresses	Protocol	Port
lc-eu2.extremecloudiq.com	3.64.95.0/29	HTTPS	443

Note
Rancher connection is required for day-to-day service operation. (It creates a tunnel to Kubernetes cluster for CloudOps remote access/management.)

For NAT deployments where you deploy your cluster with private addressing, you must provide the CloudOps team with direct admin access to the cluster nodes in your internal network. Use the mappings in the following table to map inbound ports on the public side of the NAT router to specific cluster nodes and ports in your private network.

Note
Make sure to let the CloudOps team know which IP address you are using for inbound connections. As a best practice, use the first public IP address, although you can use another address, including a public IP address that is dedicated to this connection type.

Table 15: Inbound Traffic Port Mapping (when using NAT)

Source	Service	Inbound IP (public NAT)	Inbound Port (public NAT)	Forward to UCP Node	On Port	Protocol
Restricted IP list Extreme Bastion servers: • Raleigh Bastion Host: 134.141.117.45/32 • Salem Bastion Host: 134.141.4.8/32 • San Jose: 208.185.247.165 • Thornhill: 216.123.81.194 • Bangalore AMR: 14.143.116.18 • Bangalore Bagmane: 121.244.44.28 • Bangalore Ecospace: 115.110.157.126 • LC-EU2: 3.64.95.2, 3.64.95.3 • usnh: 134.141.85.210	SSH	Your public IP address	20001	Node 1	22	TCP
			20002	Node 2	22	TCP
			20003	Node 3	22	TCP
			20004	Node 4	22	TCP
			20005	Node 5	22	TCP
			20006	Node 6	22	TCP
	UCP Remote Access	Your public IP address	20501	Node 1	5825	HTTPS
			20502	Node 2	5825	HTTPS
			20503	Node 3	5825	HTTPS
			20504	Node 4	5825	HTTPS
			20505	Node 5	5825	HTTPS
			20506	Node 6	5825	HTTPS

Note
For SSH or UCP Remote access, inbound access is needed only on-demand for the initial deployment, software upgrade, or issue troubleshooting. For <rdc>-inlets, inbound access is needed on an ongoing basis.

Access for Readiness Assessment

The Readiness Assessment tool requires access to the following address during initial deployment. The tool accesses this server to test your firewall configuration and initial settings. Run the Readiness Assessment before you install the ExtremeCloud IQ engine.

Table 16: Access for Readiness Assessment

Domain Name	IP Addresses	Protocol	Port
ucp0-console-extremecloudiq.com	18.192.120.159	HTTPS	443

Access for Production Sanity Verification

The Extreme QA team will run production santify verification after the release upgrade to make sure all of the services are still working properly. The following table shows the connection info they'll use, including the public-facing IPs from which they'll connect (column 1) and the destination ports mappings to access the cluster (column 5).

Table 17: Inbound Traffic

Source IPs	Protocol	IP Port	Description	Destination Port Mapping
Restricted IP list Extreme Bastion servers: • Raleigh Bastion Host 134.141.117.45/32 • Salem Bastion Host 134.141.4.8/32 • San Jose: 208.185.247.165 • Thornhill: 216.123.81.194 • Bangalore AMR: 14.143.116.18 • Bangalore Bagmane: 121.244.44.28 • Bangalore Ecospace: 115.110.157.126 • LC-EU: 3.64.95.7	HTTPS (TCP)	443	GDC Web Service RDC Web Service	IP1:443 → VRRP1:443 IP4:443 → VRRP4:443
	TCP	80	CAPWAP Services	IP1:80 → VRRP1:80 IP2:80 → VRRP2:80 IP3:80 → VRRP3:80 IP4:80 → VRRP4:80
	UDP	12222	CAPWAP Services	IP1:12222 → VRRP1:12222 IP2:12222 → VRRP2:12222 IP3:12222 → VRRP3:12222 IP4:12222 → VRRP4:12222
	TCP	2083	RADSEC Proxy	IP1:2083 → VRRP1:2083

Source Address Information

For installations where APs are installed off-premises and connecting for service through a firewall, relax the access rules to specific service ports because source addresses are not always deterministic.

These settings are required to support remote diagnostics and to set up validation operations.

Table 18: Source address information (examples):

Source IP	Port	Description	Action
0.0.0.0/0	TCP 80	AP CAPWAP registration	Allow
0.0.0.0/0	TCP 443	ExtremeCloud IQ login access and software updates	Allow
0.0.0.0/0	TCP 2083	RADSEC	Allow
0.0.0.0/0	UDP 12222	AP CAPWAP	Allow

Supported Hardware for Managed Orchestration

ExtremeCloud Edge - Managed Orchestration deployments of Universal Compute Platform support the following hardware appliances.

Table 19: Supported Hardware for ExtremeCloud Edge - Managed Orchestration

Hardware Appliance	Details
3160C	Ports: • 2 x 1/10 Gbps ICC Ports/RJ45 • 2 x 10/25 Gbps Data 1-2/SFP28 • 2 x 10/25/50/100 Gbps Data 3-4/QSFP For additional server specifications, along with hardware install information, see Extreme Networks Universal Compute Platform Appliance 3160C Installation Guide.
4120C-1	Ports: • 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 For additional server specifications, along with hardware install information, see Extreme Networks Universal Compute Platform Appliance 4120C Installation Guide.

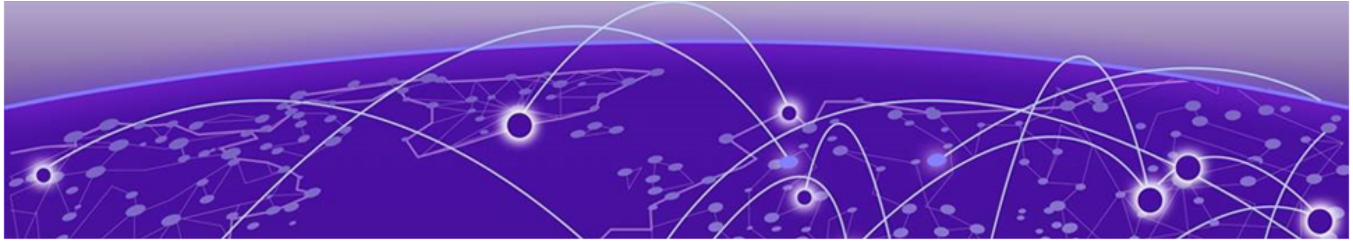

Configure the Deployment

Complete the following tasks to configure nodes for Managed Orchestration.

Table 20: Configure the Deployment

Step	Procedure	Description
1	Connect the hardware appliance to the network.	
2	Run the Basic Configuration Wizard on page 30	For each node, run the wizard to assign basic network settings.
3	Upgrade the Appliance Software on page 33	For each node, upgrade the Universal Compute Platform software.
4	Configure VRRP (VIP) on page 36	For each node, configure VRRP addresses and settings.
5	Configure the Cluster Settings on page 37	Configure the ExtremeCloud Edge - Managed Orchestration cluster.
6	Run Readiness Assessment on page 40	Run the assessment from a single node to test whether your planned configuration works.
7	Install ExtremeCloud IQ Engine on page 41	Install the engine on one cluster node.

What to do Next

Onboard Cluster to ExtremeCloud IQ on page 44—Onboard the configured cluster to the cloud and complete an Extreme CloudOps registration request. CloudOps will complete the deployment process by installing the required applications and will notify you upon completion.

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 if you want to update network settings. Settings are configured within the following groups:

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

Note
The wizard prompts you with a series of yes or 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`.

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`.

   Note
   Do not change the ICC addressing configuration after 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, the only fix is to reinstall and reconfigure

   e. Select `y` to accept the ICC configuration.

   The ICC configuration gets accepted, and the default data port configuration displays.
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.
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 screen (and in the following table). 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.

Table 21: 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

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.

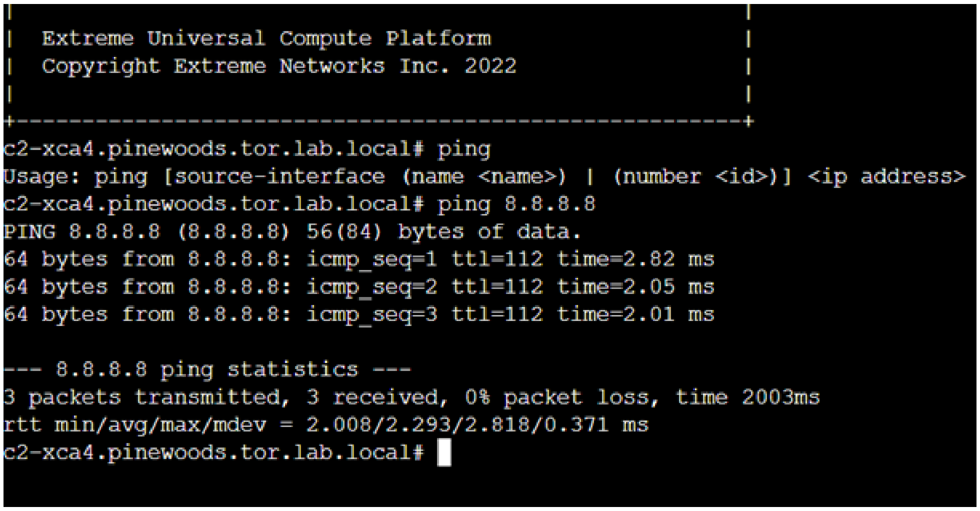

```
|                                                        |
|  Extreme Universal Compute Platform                    |
|  Copyright Extreme Networks Inc. 2022                  |
|                                                        |
+--------------------------------------------------------+
c2-xca4.pinewoods.tor.lab.local# 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 4: Example ping command

Upgrade the Appliance Software

Before configuring the cluster, use your Extreme Support account to download the latest revision of Universal Compute Platform software from the Extreme Networks Support Portal.

The image file uses one of these extensions:

- .jmx (for 3160C)
- .rcx (for 4120C)

1. Log in to the Admin user interface at: https://*node ip*:5825
2. Go to **Administration** > **System** > **Software Upgrade**.
3. Select **Upload**.

4. Upload the desired revision of Universal Compute Platform.

   Note
   A best practice is to upgrade each of the nodes on a new cluster to the latest revision before proceeding with the cluster set up and configuration.

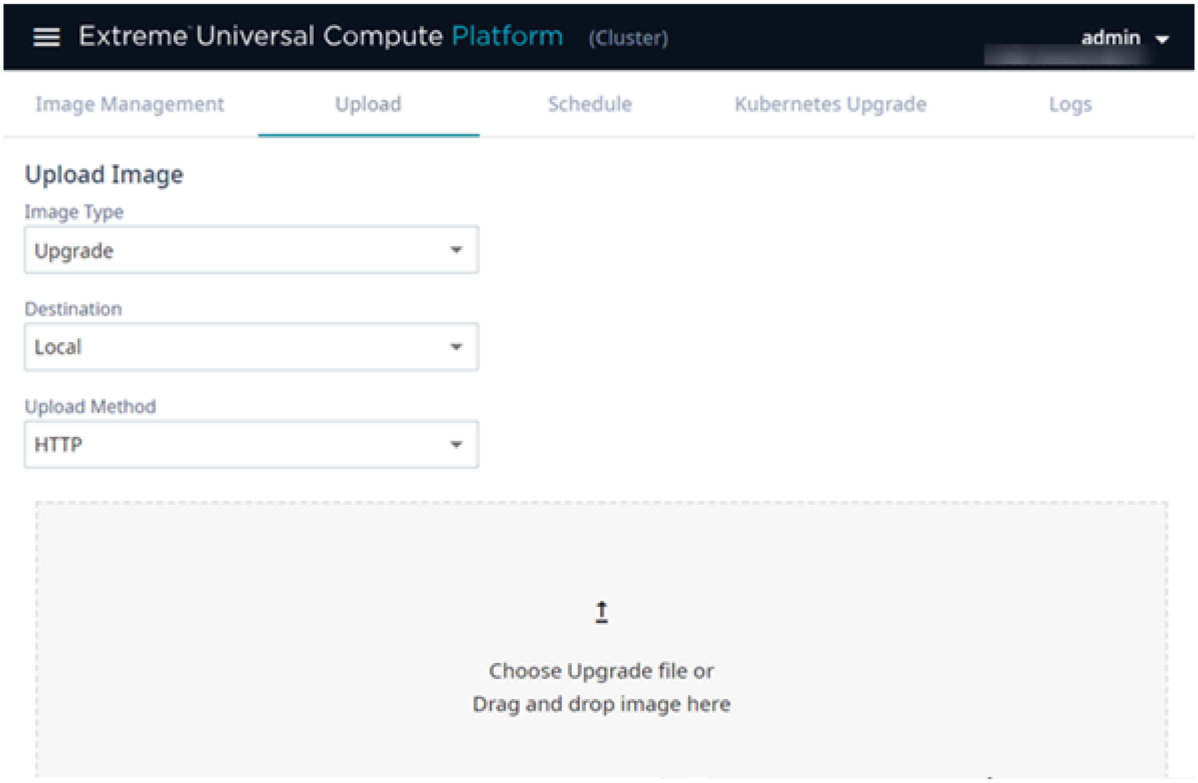

Upload Image

Image Type: Upgrade

Destination: Local

Upload Method: HTTP

Choose Upgrade file or
Drag and drop image here

Figure 5: Select the upgrade image

5. From the **Image Management** Tab, select the Upgrade image, and select ↥.
6. In the popup window, complete the following fields:
   - **Image**—Select the platform image file.
   - **Backup System Image to**—Select **Local**.
   - **Upgrade**—Select **Now**.

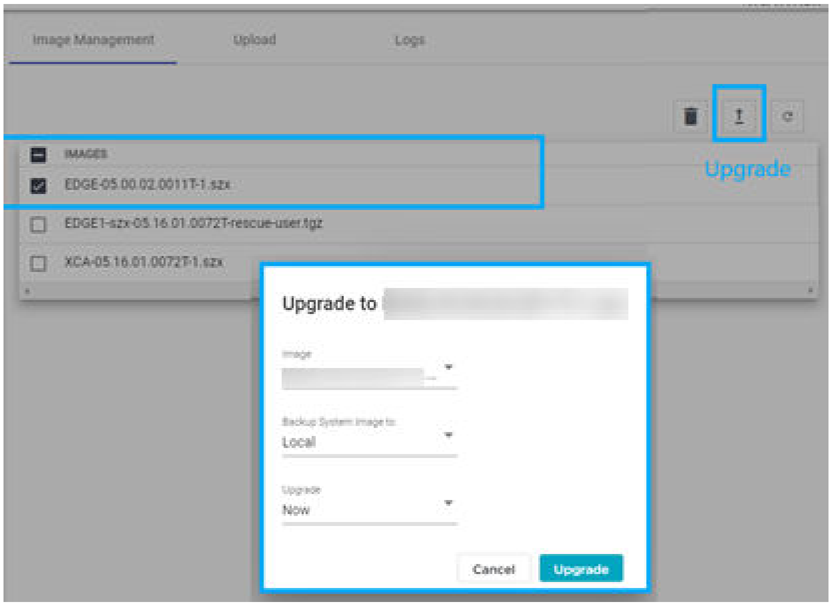

IMAGES

☑ EDGE-05.00.02.0011T-1.szx

□ EDGE1-szx-05.16.01.0072T-rescue-user.tgz

□ XCA-05.16.01.0072T-1.szx

Upgrade to

Image

Backup System Image to: Local

Upgrade: Now

Figure 6: Upgrade the selected image

7. Select **Upgrade**.

When all nodes in the cluster are upgraded to the latest revision, proceed to IP Address Configuration on page 35.

IP Address Configuration

Use the configuration wizard to initialize nodes in a cluster to the pre-determined IP addresses.

Note
IP address configuration for interfaces on the cluster must be set only once. If you change IP addresses after initial deployment (for example, due to a cluster relocation), you must rebuild and re-deploy the cluster, and re-install the application.

The following is example information that must be gathered during the prerequisite stages for each node in the cluster, and for the ICC VRRP:

IP Address

A unique IP address for each node. Example:192.227.109.81

Mask

Common Mask. Example:/26 (255.255.255.192)

Gateway

Common Gateway. Example:192.227.109.65

VRRP Precedence

Common Router ID with a unique precedence level for each node. Provide a unique precedence value for each node.

For example:

- Node1 - 100 Router ID1
- Node2 - 75 Router ID1
- Node3 - 50 Router ID1
- Node4 - 25 Router ID1
- Node5 - 10 Router ID1
- Node6 - 01 Router ID1

Configure VRRP (VIP)

Take the following steps to configure the Virtual Router Redundancy Protocol (VRRP) IP addresses.

1. Navigate to **Administration** > **System** > **Network Setup**.
2. From the **Interfaces** list, select the data access interface that you configured from the **System Startup Wizard** for (Port 1).

   The Port Configuration Settings menu displays.
3. Under **VRRP**, provide a list of IP addresses that will be offered via VRRP.
4. Set the **Priority** and **Router ID**.

   Each node must have the same VRRP IP addresses and the same Router ID, but a unique Priority setting. The Priority setting determines which node in the cluster is the Primary node. The node with the higher priority is considered the default Primary node.

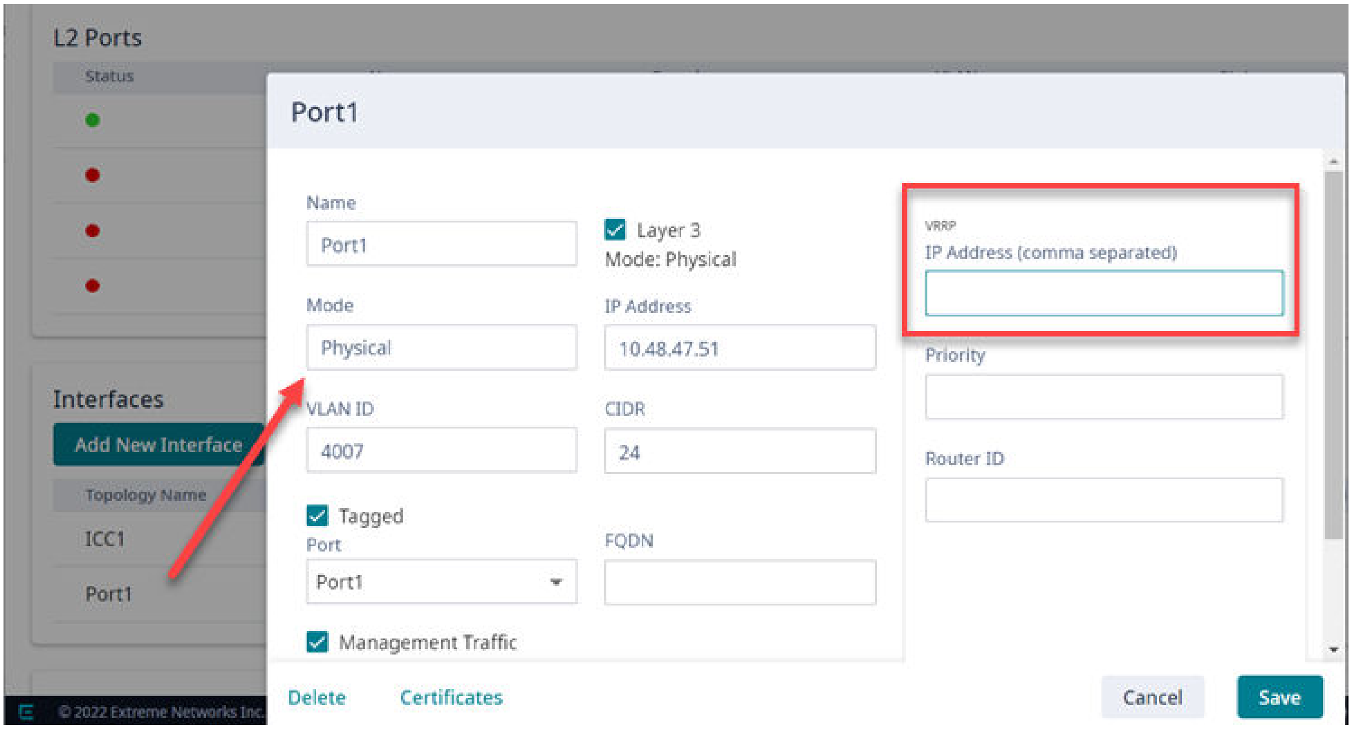

Port1

Name: Port1

☑ Layer 3

Mode: Physical

Mode: Physical

IP Address: 10.48.47.51

VLAN ID: 4007

CIDR: 24

☑ Tagged

Port: Port1

FQDN:

☑ Management Traffic

VRRP

IP Address (comma separated):

Priority:

Router ID:

Figure 7: User Interface showing properties window for Port 1

5. Repeat this process in each of the nodes of the cluster.

Configure the Cluster Settings

An engine is an instance of a containerized application. This process follows the user interface to configure the orchestration engine settings. From the management IP address, log into the user interface using the admin credentials that you configured when you ran the basic configuration wizard.

To configure the cluster, go to **Cluster Settings** > **Cluster Configuration** and use the wizard to configure the cluster using the order shown on screen:

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

Complete these steps to configure the cluster:

1. From the **Deployment Type** drop-down list, select **ExtremeCloud Edge - Managed Orchestration** and select **Next**.
2. In the **Cluster Mode** section, select **Cluster** mode and do the following:
   a. From the **Number of Nodes to Add** drop-down, select the number of nodes for the cluster.
   b. From the **Number of Availability Zones** drop-down, select the number of availability zones for the cluster. You can assign up to three zones (the default is one).
   c. Enter the ICC IP addresses that you have assigned to each node.

   Note
   If you are deploying multiple availability zones, make sure that each ICC address falls under the zone where you want to deploy that node. Note that each zone requires a minimum of three nodes. For more information, see Availability Zones on page 18.

   d. Select **Next**.

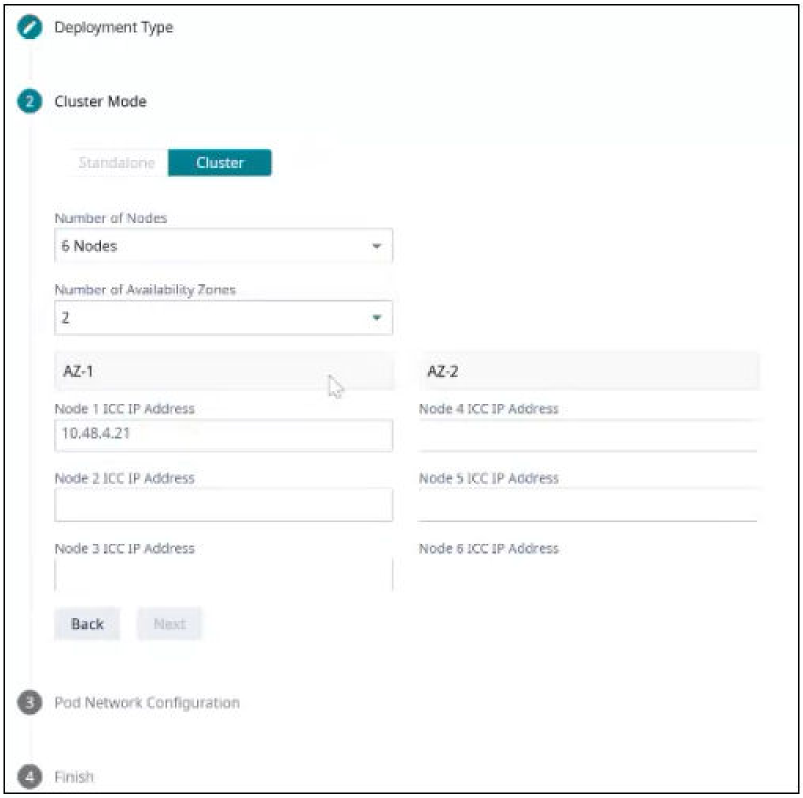

Deployment Type

2 Cluster Mode

Standalone Cluster

Number of Nodes
6 Nodes

Number of Availability Zones
2

AZ-1

Node 1 ICC IP Address
10.48.4.21

Node 2 ICC IP Address

Node 3 ICC IP Address

AZ-2

Node 4 ICC IP Address

Node 5 ICC IP Address

Node 6 ICC IP Address

3 Pod Network Configuration

4 Finish

Figure 8: Appliance Configuration

3. Provide the following network settings for **Pod Network Configuration**:
   - Pod Network IP Address—Enter a network address. The default is 10.96.0.0.
   - Pod Network CIDR—Enter the number of digits for the pod network mask. The default is 16.
   - Service Network IP Address—Enter a network address. The default is 10.97.0.0.
   - Service Network CIDR—Enter the number of digits for the service network mask. The default is 16.

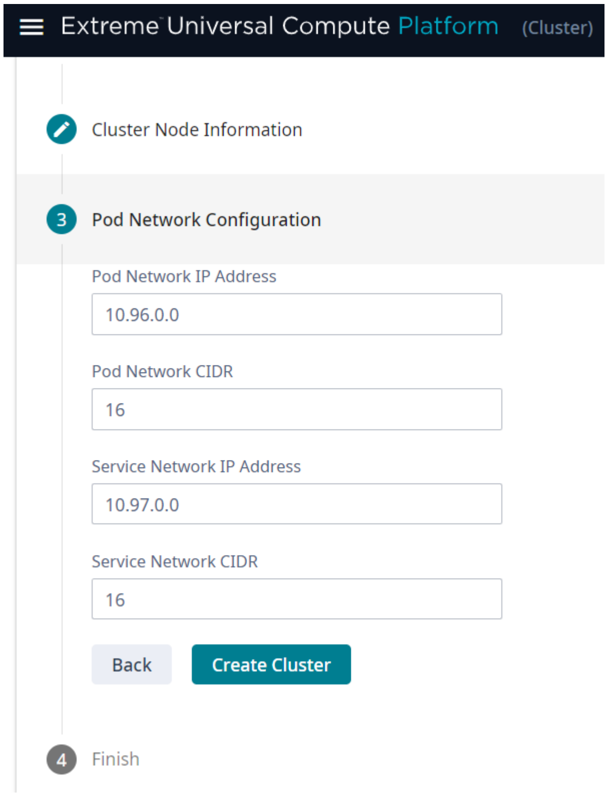

Extreme Universal Compute Platform (Cluster)
Cluster Node Information
3 Pod Network Configuration

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
4 Finish

Figure 9: Pod Settings

4. Select **Create Cluster**.

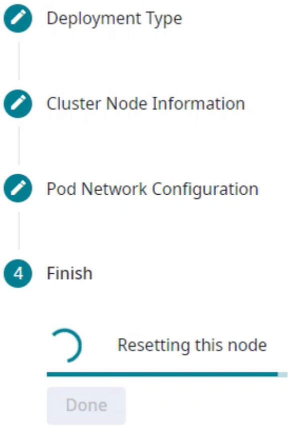

5. Select **Done**.

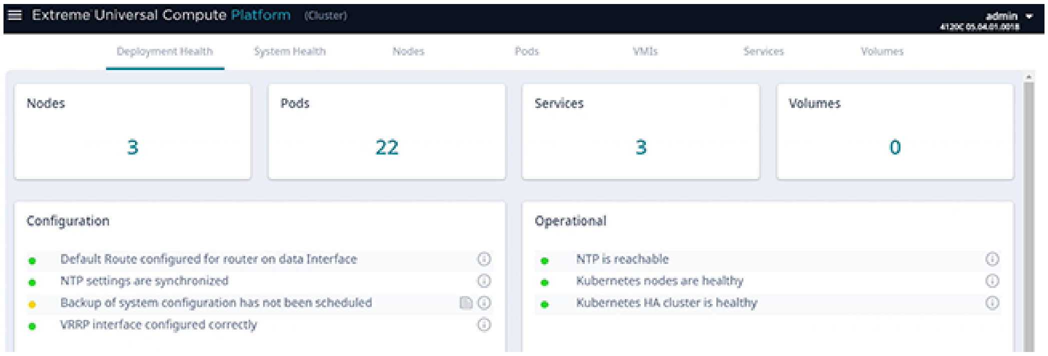

Nodes 3

Pods 22

Services 3

Volumes 0

Configuration

- ● Default Route configured for router on data Interface
- ● NTP settings are synchronized
- ● Backup of system configuration has not been scheduled
- ● VRRP interface configured correctly

Operational

- ● NTP is reachable
- ● Kubernetes nodes are healthy
- ● Kubernetes HA cluster is healthy

> Note
> Once the cluster has been created, going to the **Cluster Configuration** page displays a read-only view of cluster configuration settings.

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 from a single node 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.

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 22: 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.
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

ExtremeCloud™ IQ is the only available engine for an ExtremeCloud Edge - Managed Orchestration deployment. Install the ExtremeCloud IQ engine once from a single node.

> **Note**
> Installing the ExtremeCloud IQ engine prepares the cluster to receive the ExtremeCloud IQ application but does not install the application. Extreme CloudOps completes the application installation after you onboard the cluster and register your account.

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

   The system is prepared to receive the ExtremeCloud IQ application(s). Please proceed to onboard and register the cluster.

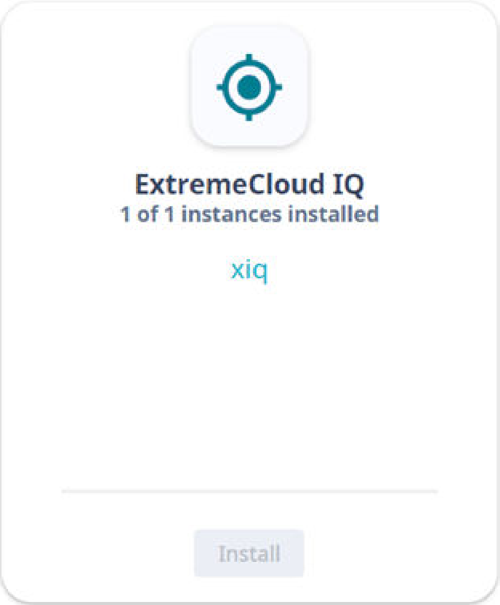

ExtremeCloud IQ
1 of 1 instances installed
xiq

Figure 10: Installed ExtremeCloud IQ Engine Instance

Note
An Edge Cloud deployment of ExtremeCloud IQ must be configured in a cluster of three or more nodes in multiples of three. The minium number of cluster nodes is three for ExtremeCloud IQ only, and six nodes if you're also deploying other applications. ExtremeCloud IQ is not supported in stand-alone mode, requires a cluster, and does not support engine types other than ExtremeCloud IQ.

Network Service Configuration

Map each core service set to the corresponding Virtual Router Redundancy Protocol (VRRP). Assign a VRRP virtual router address for each set of services. VRRP enables a virtual router to act as the default network gateway, improving host network reliability and performance.

Validate the Cluster

Click the **Deployment Health** tab for information.

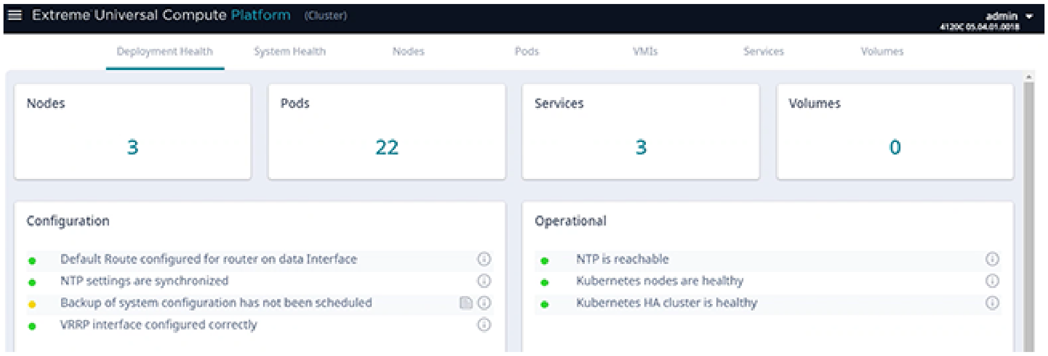

Nodes 3

Pods 22

Services 3

Volumes 0

Configuration

- ● Default Route configured for router on data Interface
- ● NTP settings are synchronized
- ● Backup of system configuration has not been scheduled
- ● VRRP interface configured correctly

Operational

- ● NTP is reachable
- ● Kubernetes nodes are healthy
- ● Kubernetes HA cluster is healthy

Onboard Cluster 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.
2. Initiate action for the ExtremeCloud IQ Operations team to deploy a Regional Data Center (RDC) for the cluster.
3. Register your ExtremeCloud IQ account.
4. Onboard your devices and operate the account.

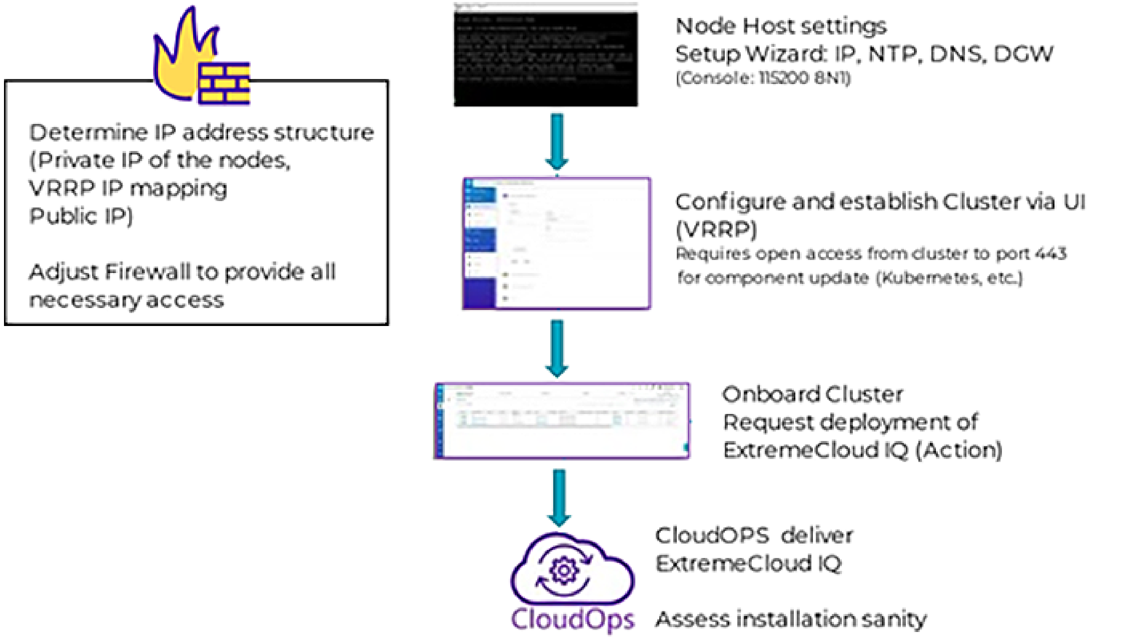
Determine IP address structure
(Private IP of the nodes,
VRRP IP mapping
Public IP)
Adjust Firewall to provide all necessary access
Node Host settings
Setup Wizard: IP, NTP, DNS, DGW
(Console: 115200 8N1)
Configure and establish Cluster via UI (VRRP)
Requires open access from cluster to port 443 for component update (Kubernetes, etc.)
Onboard Cluster
Request deployment of ExtremeCloud IQ (Action)
CloudOPS deliver ExtremeCloud IQ
Assess installation sanity
CloudOps

Figure 11: ExtremeCloud Edge Deployment Workflow

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

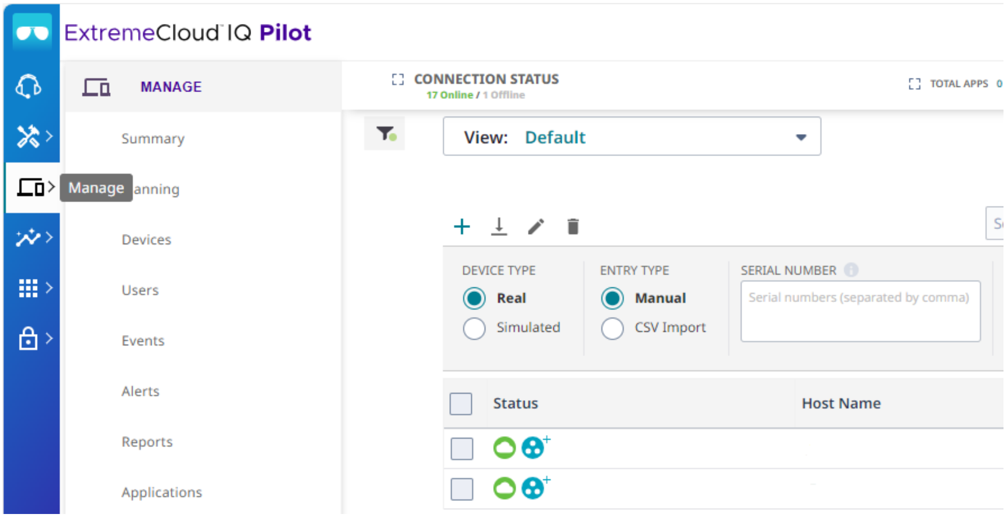

Figure 12: Add Cluster to ExtremeCloud IQ

The **Device Make** field displays.

4. From the **Device Make** menu, select **Universal Appliance**.
5. 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.

6. Select **Actions** > **Applications** > **ExtremeCloud IQ Manage**.

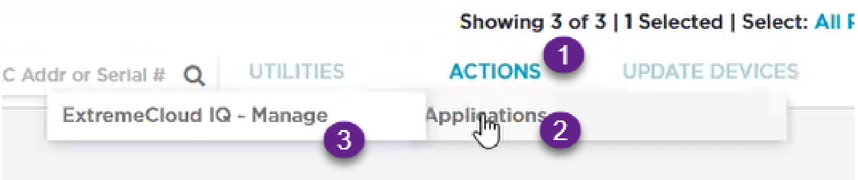

Figure 13: ExtremeCloud IQ Actions menu

This initiates the action for ExtremeCloud IQ OPs to deploy a Regional Data Center (RDC) for the cluster.

7. Fill out the online form:

   > Note
   > Required fields are noted with an asterisk.

   - Customer Information
   - Primary Technical Contact
   - Secondary Technical Contact
   - Notification List — Provide a list of email addresses for notification.
   - Nightly Backup
   - Scheduled Upgrades
   - RDC Name — Provide a meaningful name, up to 6 characters. The system will verify that the name is available.
   - IP Address Mapping — Provide the mapping between the external Public IP Address to the internal virtual VRRP IP Address for each service set.

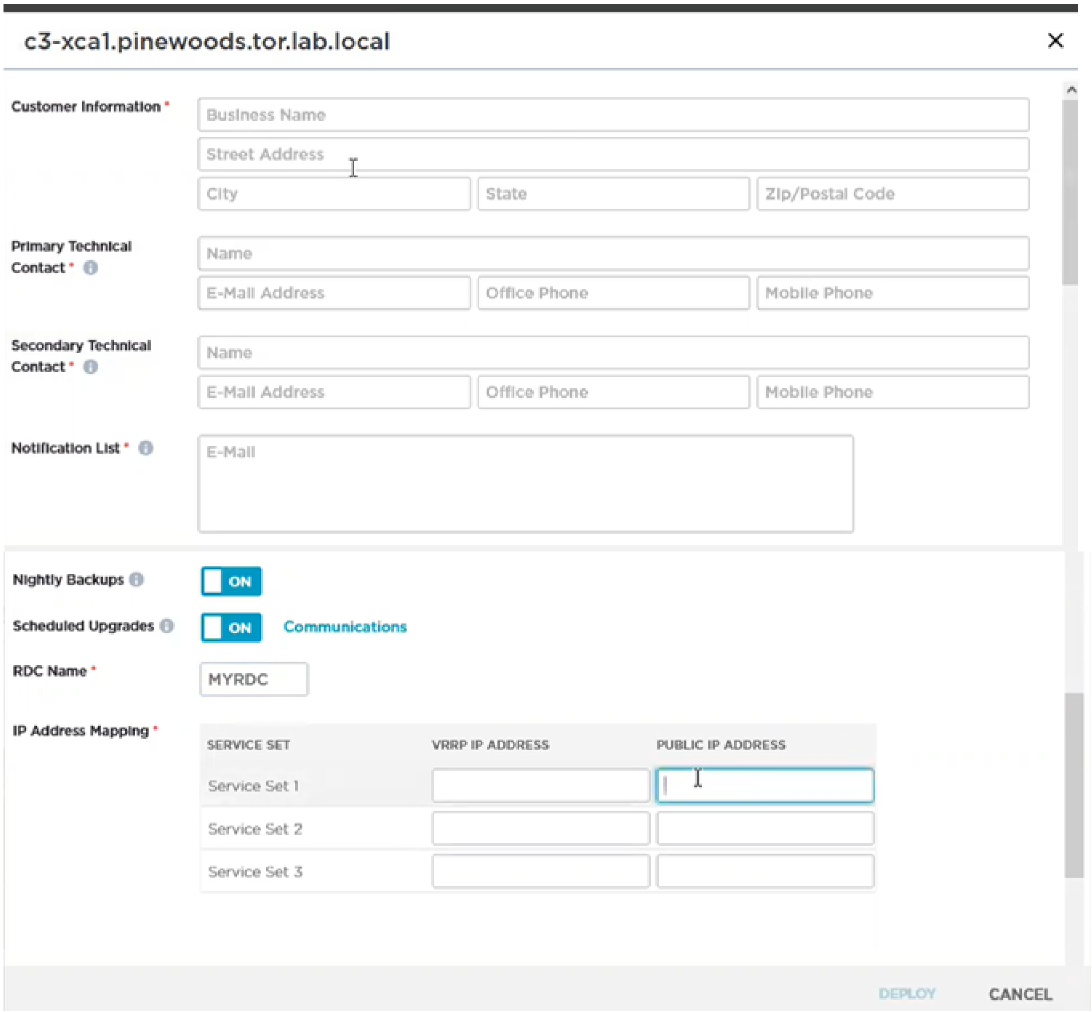

c3-xca1.pinewoods.tor.lab.local ×

Customer Information *: Business Name / Street Address / City / State / Zip/Postal Code

Primary Technical Contact *: Name / E-Mail Address / Office Phone / Mobile Phone

Secondary Technical Contact *: Name / E-Mail Address / Office Phone / Mobile Phone

Notification List *: E-Mail

Nightly Backups: ON

Scheduled Upgrades: ON Communications

RDC Name *: MYRDC

IP Address Mapping *

SERVICE SET	VRRP IP ADDRESS	PUBLIC IP ADDRESS
Service Set 1		
Service Set 2		
Service Set 3		

Figure 14: ExtremeCloud IQ Deploy a Cluster Form

8. Select **Deploy**.

   A ticket is generated for ExtremeCloud IQ OPs. Operations personnel will provide an estimate for the expected deployment schedule.

   During deployment, the OPs team will do the following:

   - Deploy ExtremeCloud IQ software to the on-premise hosts
   - Validate the deployment to ensure that the site is deployed and operating correctly
   - Once validated, OPs will provide notification of readiness
   - Provide the installation token that enables customers to create accounts directly on the newly deployed ExtremeCloud IQ private Regional Data Center (RDC).

   Note
   You can view the status of the deployment process from the **Application Status** column on the **Device List**

Account Registration

For information about creating accounts after you set up ExtremeCloud Edge, consult the Managed Service Partner (MSP) documentation.

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.

Configure Persistent Connection to ExtremeCloud

Use this optional procedure to configure a persistent connection to the ExtremeCloud network using the **CloudOps Management** setting. CloudOps uses this persistent connection to manage software delivery for the cluster.

By default, the setting is On for all Managed Orchestration clusters. By turning the setting to Off, you can limit software installs and upgrades to specific maintenance windows.

1. Go to **Engines** > **Installation**.
2. Select the **ExtremeCloud IQ** engine.
3. Select the **Settings** tab.

4. Set **CloudOps Management** to one of the following settings:
    - On—The cluster maintains a persistent connection to the ExtremeCloud network. This is the default setting.
    - Off—The cluster disconnects from the ExtremeCloud network. A customer administrator or CloudOps administrator must turn this setting back to On before CloudOps can install or upgrade software.
5. Select **Save**.

Note
If **CloudOps Management** is Off, the setting can be turned back on by an administrator at the customer organization or by the CloudOps team (using remote GUI access).

Appendix

Appendix A: ExtremeCloud IQ - Site Engine Integration with ExtremeCloud Edge

To integrate ExtremeCloud IQ - Site Engine with an ExtremeCloud Edge - Managed Orchestration cluster, you must edit a pair of parameters in the `NSJBoss.properties` file on the ExtremeCloud IQ - Site Engine server. By default, these parameters point to ExtremeCloud IQ in the public cloud. However, you must edit the file so that these parameters point to services on the local ExtremeCloud Edge - Managed Orchestration cluster.

- `extreme.xiq.baseUrl`—The Base URL is used for authentication and points to the GDC (Global Data Center).
- `extreme.xiq.redirectorurl`—The Redirector URL is a public URL that redirects the application to the local RDC (Regional Data Center).

To configure the integration:

1. On the ExtremeCloud IQ - Site Engine server, open the `NSJBoss.properties` file and edit the following parameters:
    - Set `extreme.xiq.baseUrl=https://<EDGE_name>-g1.extremecloudiq.com` where `<EDGE-name>` is the name of your installation (aka prefix name).
    - Set `extreme.xiq.redirectorurl=https://hac.extremecloudiq.com`
2. Run the command `systemctl restart nsserver` to restart the server with the new settings.

Appendix B: Migrate Virtual IQ Account

This Appendix describes how to migrate a Virtual IQ (VIQ) account to a new Regional Data Center (RDC). To migrate the account, complete each of the subsequent procedures in order:

1. Export VIQ Account

2. Import VIQ Account

> Note
> Moving the VIQ account also moves the account inventory (for example, devices, floor plans, private pre-shared keys) as well as configurations and assignments.

Export VIQ Account

Use this procedure to create and download an export file for a VIQ account.

1. In ExtremeCloud IQ Pilot, go to **Global Settings** and select **VIQ Management.**
2. Create a backup of the current VHM:
   a. Under **VIQ Management**, select **BACK UP NOW**.
   b. Select **YES**.

      VIQ suspends itself until the backup completes.
3. Export the VHM to a local drive:
   a. Go to **Global Settings** and select **VIQ Management**.
   b. Select **Export VIQ**.
   c. In the **VIQ Export** popup window, select **Export Now**.
   d. Click **YES**. VIQ suspends itself until the Export completes.
   e. Once the export completes successfully, select **OK**.

      > Note
      > If the export fails, click the **Detailed Report** link to get a detailed report on the issue.

What to do Next

After the export file downloads, you can import the file into a different Regional Data Center (RDC).

Import VIQ Account

Use this procedure to import the VIQ export file into the new RDC. Note the following:

- If a conflict occurs, imported objects get renamed.
- Source and destination VHMs must be the same version. Otherwise, an incompatible data scheme occurs.

1. From ExtremeCloud IQ Pilot, go to **Global Settings** > **VIQ Management**.
2. Create a backup of the current VHM:
   a. Under **VIQ Management**, select **BACK UP NOW**.
   b. Select **YES**.

      VIQ suspends itself until the backup completes.

3. Import the VHM export file that you created in the preceding procedure:
    a. Select **Import VIQ**.
    b. Select **Import VIQ from ExtremeCloud IQ**.
    c. Select **Choose** and then browse and select the VHM export file.
    d. Select **Import Now**.
    e. After the import completes, select **OK**.

    Note
    - If the import fails, download the log file for information on the issues.
    - If you need to roll back the import, restore the backup.

Appendix C: Replace or Add a Node

Use the procedures in this Appendix if you need to replace a node or add a node.

Prepare to Replace a Node

1. Gather the IP address settings of the failed node.

    Unless stated otherwise, you will set the new node with the same IP address values as the unit being replaced:
    - ICC Interface IP Address—For the ICC interface, you must assign a new IP address to the replacement node.
    - Data Port Interface IP Address
    - DNS Server Address
    - NTP Server Address
2. Configure the VRRP priority for the replacement node.

    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.

3. Use the Basic Configuration Wizard to configure the replacement unit.

    This is required if you are replacing the unit hardware. Node Replacement initially resets the node connections. It may not require new hardware.

    For information about the Basic Configuration Wizard, see the appropriate Deployment Guide.
4. Upgrade the Appliance Software on page 33 and upgrade the node to the current software version.

What to do Next

After you have gathered the necessary information and verified the software version of all nodes in the cluster, go to the Replace Node on page 52 procedure.

Replace Node

Replacing a node in a cluster is performed when a node has failed and must be replaced. The replacement node gets delivered in a reset state. After initializing the node for its network presence, the new node is added to the cluster and assumes the service load of the removed node.

Note
Before you replace a node, review the information in Prepare to Replace a Node on page 51.

From the primary node in the cluster (Node 1), take the following steps:

1. Go to **Cluster Settings** > **Node Replacement**.

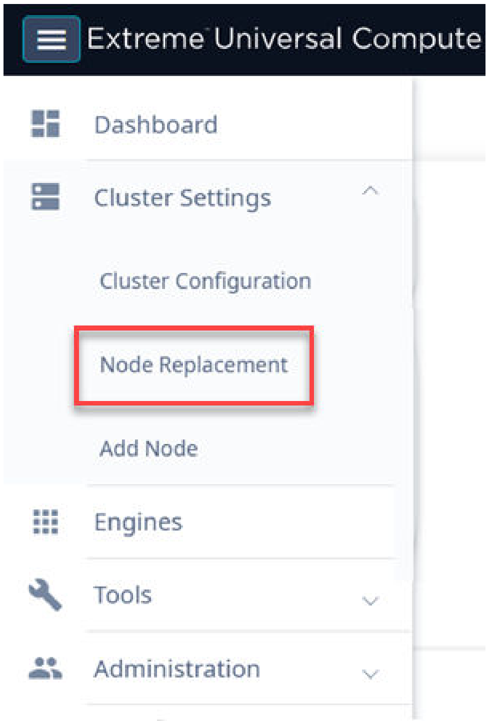

Figure 15: Node Replacement

2. Select the failed node and select **Next**.

   Existing credentials are used to establish connection to the failed node. Configuration and services information is transferred from the primary node to the failed node in an effort to re-establish a connection.

   If it is necessary to replace the node hardware, refer to the Installation Guide for your Universal Compute Platform Appliance model for detailed information.

Add Node

Note

Before adding a new node, you must configure the new node and ensure that it is running the current software version. Refer to Prepare to Replace a Node on page 51.

A node is one appliance. Universal Compute Platform multi-node clusters can be deployed with three or more nodes, with the ability to scale up when the cluster reaches capacity. With a multi-node cluster, the cluster size must be a multiple of three (for example, three, six, and nine are acceptable cluster sizes).

To add one or more nodes to a cluster, take the following steps:

1. Go to **Cluster Settings** > **Add Node**.
2. In the **Number of Nodes to Add** drop-down, select the number of nodes that you are adding.
3. 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 you have multiple zones configured already within the cluster. The new nodes that you want to add must be added to the configured zones.

   Note

   When adding new nodes to an existing cluster that is split into multiple availability zones, keep the following points in mind:

   - The cluster configuration, after the nodes are added, must align with availability zones requirements for an equal number of nodes per zone.
   - The cluster size, after the nodes are added, must be a multiple of three.
   - There is no option to reconfigure the number of zones after the initial cluster is created. This limitation also applies when adding nodes to an existing cluster.

   For example, when adding nodes to a six-node and two zone cluster, you can add nodes only in multiples of six or the node addition would violate one of these rules. However, when adding nodes to a nine-node cluster that has three zones, you have the option to add three nodes (one node per zone) giving you a twelve-node cluster with three zones of four nodes each.

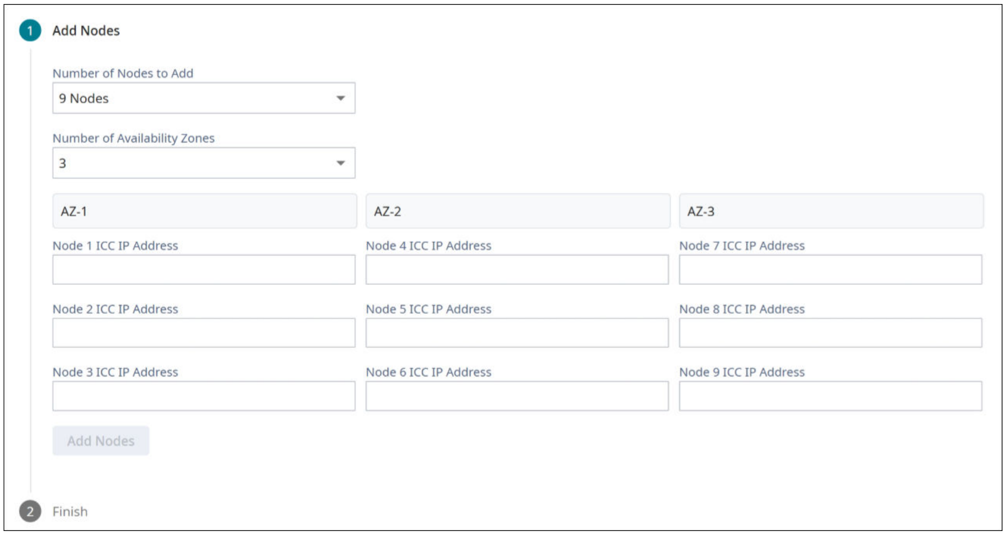

1 Add Nodes

Number of Nodes to Add: 9 Nodes

Number of Availability Zones: 3

Figure 16: Add Nodes Example with Availability Zones

4. Enter the ICC IP Address for each node and then select **Add Nodes**.
5. Select **OK** to begin the Add Node process.

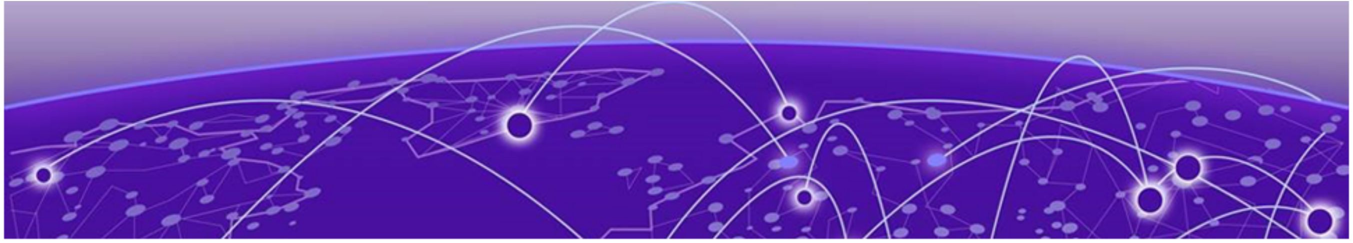

Index