



Cisco Intercloud Fabric Provider Platform Installation Guide, Release 2.3.1

First Published: November 13, 2015

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2015 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1

Overview 1

Cisco Intercloud Fabric Provider Platform 1

CHAPTER 2

Cisco ICFPP Deployment Options 3

Deployment Options 3

Standalone Configuration 3

Cluster Configuration 4

Deployment Workflows 5

CHAPTER 3

Installation Requirements 7

System Requirements 7

Hypervisor Requirements 8

Port Requirements 8

Information Required for Configuration and Installation 9

CHAPTER 4

Installing Cisco ICFPP on VMware 13

Cisco ICFPP Software 13

Workflow for VMware Environments 14

Installing Cisco ICFPP on VMware 14

Configuring the IP Address for Network Access 17

CHAPTER 5

Installing Cisco ICFPP on OpenStack 19

Workflow for OpenStack Environments 19

Installing Cisco ICFPP on OpenStack 19

Configuring Cisco ICFPP for Cisco Intercloud Fabric Director 21

CHAPTER 6

Upgrading Cisco ICFPP 25

Upgrading Standalone Nodes or Multiple-Node Clusters 25

Supported Upgrade Paths	25
Restarting Services Automatically	25
Upgrading a Standalone Node	26
Upgrading a Multiple-Node Cluster	27

CHAPTER 7

Configuring Cisco ICFPP for Clusters	31
Workflow for Configuring Clusters	31
Configuring a Primary Node	32
Configuring a Service Node	33
Configuring Additional Storage	34
Configuring NFS	34
Configuring a Cinder Volume	35
Configuring HA	37
Configuring VIP Access for HA in OpenStack	39
Moving from a Standalone Setup to a Cluster	42
Restoring a Database onto an Existing HA Pair	43
Monitoring HA Status	44
Viewing HA Syslog Messages	45

CHAPTER 8

Configuring VMware vCloud Director for Cisco ICFPP	47
Configuring VMware vCloud Director	47
Workflow for Integrating VCD with Cisco ICFPP	50
Creating an External Network	51
Adding a vShield Edge Gateway on an Org VDC	52
Creating an Org VDC Internal Network	53
Creating a Catalog	55
Verifying NAT and Firewall Service Configuration	56

CHAPTER 9

Additional Information	59
Related Documentation for Cisco Intercloud Fabric Provider Platform	59
Obtaining Documentation and Submitting a Service Request	60
Documentation Feedback	60



Overview

- [Cisco Intercloud Fabric Provider Platform, page 1](#)

Cisco Intercloud Fabric Provider Platform

Cisco Intercloud Fabric Provider Platform (ICFPP) simplifies the complexity involved in working with a variety of public cloud APIs, and enables cloud API support for service providers who currently do not have API support. Cisco ICFPP provides an extensible adapter framework that allows integration with a variety of provider cloud infrastructure management platforms, such as OpenStack, CloudStack, VMware vCloud Director, and any other API that can be integrated through a software development kit (SDK) provided by Cisco.

Currently, service providers have their own proprietary cloud APIs—such as Amazon EC2 and Windows Azure—that give customers limited choices and do not provide an easy method for moving from one provider to another. Cisco ICFPP abstracts this complexity and translates Cisco Intercloud Fabric cloud API calls to cloud platform APIs of different provider infrastructure platforms, giving customers the option of moving their workloads regardless of the cloud API used by the service provider.

Many service providers do not provide cloud APIs that Cisco Intercloud Fabric can use to deploy customers' workloads. One option for these providers is to provide direct access to the virtual machine (VM) manager's SDK or API, such as vCenter or System Center. However, this option exposes the provider environment and is not preferred by service providers because of security concerns. Cisco ICFPP, as the first point of authentication for the customer cloud when requesting cloud resources, enforces highly secure access to the provider environment. In addition, Cisco ICFPP provides the cloud APIs that are required for service providers to be part of the provider ecosystem for Cisco Intercloud Fabric.

As the interface between the Cisco Intercloud Fabric from customer cloud environments and provider clouds (public and virtual private clouds), Cisco ICFPP provides the following benefits:

- Standardizes and brings uniformity to cloud APIs, making it easier for Cisco Intercloud Fabric to consume cloud services from service providers that are a part of the Cisco Intercloud Fabric ecosystem.
- Helps secure access to a service provider's underlying cloud platform.
- Limits the utilization rate per customer or tenant environment.
- Provides northbound APIs for service providers for integration with existing management platforms.
- Supports multitenancy.

- Monitors resource usage for each tenant.
- Meters resource usage for each tenant.



Cisco ICFPP Deployment Options

- [Deployment Options, page 3](#)
- [Standalone Configuration, page 3](#)
- [Cluster Configuration, page 4](#)
- [Deployment Workflows, page 5](#)

Deployment Options

You can deploy Cisco ICFPP in the service provider data center in the following configurations:

- Standalone—Deployment on a single node.
- Multiple-node cluster—Deployment on multiple nodes including a high-availability (HA) pair and additional service nodes.

Cluster deployments are most effective when they are configured behind a load balancer. After these configurations are deployed, a provider-supplied load balancer is expected to manage cookie-based sessions and direct requests and responses appropriately.

The following topics describe these configuration options in more detail.

Standalone Configuration

In a standalone configuration, Cisco ICFPP is deployed as a single virtual appliance that provides services and acts independently of other Cisco ICFPP nodes. A standalone configuration is appropriate for environments in which redundancy is not a concern.

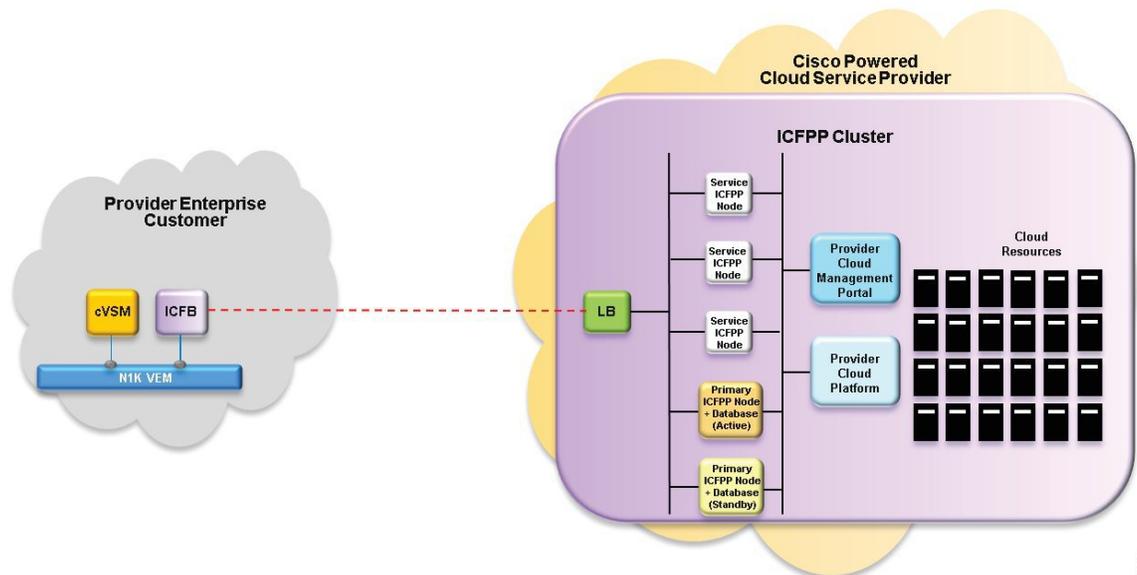
Cisco ICFPP is installed in OpenStack environments using the Standalone role only. After installation, you can configure the Cisco ICFPP virtual appliance as a primary node or service node as appropriate for your environment.

Cluster Configuration

In a cluster configuration, Cisco ICFPP supports large-scale operations in provider environments by deploying multiple Cisco ICFPP nodes. A provider-supplied load balancer distributes the load across the service nodes.

In a cluster configuration, Cisco ICFPP is deployed on multiple virtual appliances using the Primary Node and Service Node roles, as shown in the following figure.

Figure 1: Cisco ICFPP in a Multiple-Node Cluster Configuration



A multiple-node cluster contains the following components:

- Two Cisco ICFPP virtual appliances with the Primary Node role that are deployed in an HA configuration.
- Two or more Cisco ICFPP virtual appliances that are deployed as service nodes.
- A load balancer that forwards incoming traffic only to the service nodes.

HA Configuration in a Cluster

In an HA configuration, Cisco ICFPP is deployed on two virtual appliances, both using the Primary Node role. Each virtual appliance in an HA pair includes a database for replication purposes. After both Cisco ICFPP appliances are deployed, you specify which node is active and which node is standby.

The following concepts apply when Cisco ICFPP is deployed in an HA configuration:

- The Cisco ICFPP appliances in an HA pair have different management IP addresses.
- A single virtual IP address (VIP) is assigned to the active node.
- If the active node fails, the VIP is transferred to the standby node.
- When the original active node recovers, the VIP remains with the original standby node until that node fails.

Database replication works as follows:

- The active and standby nodes replicate each other's databases.
- At any time, only the database on the node with the VIP is used.
- When the database is updated on one node, the changes are replicated on the other node.

Communications in a Cluster

In a cluster configuration, Cisco ICFPP manages communications as follows:

- Each service node establishes a connection with the database on the active node in the HA pair by using the virtual IP address for the HA pair.
- The load balancer accepts requests from Cisco Intercloud Fabric Director.
- The load balancer distributes the requests to the service nodes using a round-robin algorithm.
- Each new user session is directed to a different service node.
- Subsequent requests from the same session are sent to the same service node.
- The service node responds via the load balancer.

Session Persistence

Session persistence is managed by means of a PERSISTICFPP cookie that Cisco ICFPP issues. The cookie is generated when a user logs in and it ensures that all requests from that user session are directed to the same node. If a service node fails, the load balancer forwards requests for that service node to a different service node. The new receiving node first requires Cisco Intercloud Fabric Director to log in and then accepts new requests.



Note

The service provider load balancer must be configured to persist sessions based on the PERSISTICFPP cookie.

Deployment Workflows

The deployment workflow that you use depends on whether Cisco ICFPP is deployed on VMware or OpenStack. The following table describes the high-level tasks required to deploy Cisco ICFPP in a multiple-node cluster in a VMware environment.

Table 1: Configuration Workflow for a Multiple-Node Cluster on VMware

Step	Task	Related Information
1.	Install two Cisco ICFPP virtual appliances using the Primary Node role.	Installing Cisco ICFPP on VMware, on page 14
2.	Install two or more Cisco ICFPP virtual appliances using the Service Node role.	Installing Cisco ICFPP on VMware, on page 14

Step	Task	Related Information
3.	Configure additional storage.	Configuring NFS, on page 34
4.	Configure HA on the appliances with the Primary Node role.	Configuring HA, on page 37
5.	Configure a load balancer for all service nodes in the cluster. Note The load balancer must be configured to persist sessions based on the PERSISTICFPP cookie that Cisco ICFPP issues.	Your load balancer documentation
6.	Configure communications for the cluster with Cisco Intercloud Fabric Director.	<i>Cisco Intercloud Fabric Getting Started Guide</i>

The following table describes the high-level tasks required to deploy Cisco ICFPP in a multiple-node cluster in an OpenStack environment.

Table 2: Configuration Workflow for a Multiple-Node Cluster on OpenStack

Step	Task	Related Information
1.	Install four or more Cisco ICFPP virtual appliances using the Standalone Node role.	Installing Cisco ICFPP on OpenStack, on page 19
2.	Configure two appliances with the Primary Node role.	Configuring a Primary Node, on page 32
3.	Configure the remaining appliances with the Service Node role.	Configuring a Service Node, on page 33
4.	Configure additional storage.	Configuring a Cinder Volume, on page 35
5.	Configure HA on the appliances with the Primary Node role.	Configuring HA, on page 37
6.	Configure the HA nodes to permit network traffic via the VIP address.	Configuring VIP Access for HA in OpenStack, on page 39
7.	Configure a load balancer for the service nodes in the cluster. Note The load balancer must be configured to persist sessions based on the PERSISTICFPP cookie that Cisco ICFPP issues.	Your load balancer documentation
8.	Configure communications for the cluster with Cisco Intercloud Fabric Director.	<i>Cisco Intercloud Fabric Getting Started Guide</i>



Installation Requirements

- [System Requirements, page 7](#)
- [Hypervisor Requirements, page 8](#)
- [Port Requirements, page 8](#)
- [Information Required for Configuration and Installation, page 9](#)

System Requirements

You can deploy a Cisco ICFPP virtual appliance on a system that meets the following requirements:

Requirement	Description
Four Virtual CPUs	1.8 GHz
Memory	8 GB RAM
Disk Space	Disk space that is configured as follows: <ul style="list-style-type: none">• Disk 1—100 GB for Cisco ICFPP.• Disk 2—As much memory as required to support concurrent virtual machines being moved to the provider cloud. <p>Note If additional storage is not configured, Cisco ICFPP stores VM images uploaded from Cisco Intercloud Fabric Director on the local disk. For more information on configuring additional storage, see Configuring Additional Storage, on page 34.</p>
One vNIC	Management network interface

Hypervisor Requirements

Cisco ICFPP is a virtual appliance that can be deployed on VMware vSphere Client or OpenStack KVM Hypervisor.

Hypervisor	Version
VMware	
VMware vSphere Client	5.1 (including update 1) and 5.5
OpenStack	
Cisco Intercloud Services OpenStack	Icehouse
Red Hat Enterprise Linux OpenStack Platform	Icehouse

Port Requirements

Ports must be configured as described in the following tables to ensure that Cisco ICFPP can communicate effectively on the internal private network and the public network (Internet).

Table 3: Public Internet Inbound

Protocol	Port	Allow/Deny	Description
TCP	443	ALLOW	Allows inbound HTTPS traffic from the Internet so that Cisco Intercloud Fabric for Business can reach Cisco ICFPP.

Table 4: Public Internet Outbound

Protocol	Port	Allow/Deny	Description
All	All	DENY	Cisco ICFPP does not need to send outbound traffic to the Internet.

Table 5: Internal Network Inbound

Protocol	Port	Allow/Deny	Description
TCP	443	ALLOW	Allows inbound HTTPS traffic from the internal network, so that the Cisco ICFPP web-based GUI can be accessed.
TCP	22	ALLOW	Allows inbound SSH traffic from the internal network for Cisco ICFPP administration.

Protocol	Port	Allow/Deny	Description
TCP	3306	ALLOW	Allows inbound MySQL traffic from the internal network. Required if Cisco ICFPP is configured in a multiple-node cluster.
TCP	8080	ALLOW	Allows inbound HTTP traffic for template uploads to CloudStack. Required if using the CloudStack adapter.

**Note**

To ensure that the destination systems receive communications from Cisco ICFPP, the ports in the following table must be open on any firewalls on the internal network between Cisco ICFPP and the destination systems.

Table 6: Internal Network Outbound

Protocol	Firewall Port	Allow/Deny	Description
TCP	443	ALLOW	Allows HTTPS traffic to the internal network. Required to reach the cloud provider API/SDK gateway if it is running on HTTPS.
TCP	80	ALLOW	Allows HTTP traffic to the internal network. Required to reach the cloud provider API/SDK gateway if it is running on HTTP.
TCP	3306	ALLOW	Allows outbound MySQL traffic to other Cisco ICFPP nodes on the internal network. Required if Cisco ICFPP is configured in a multiple-node cluster.
TCP/UDP	514	ALLOW	Allows syslog traffic from Cisco ICFPP to the syslog server.

Information Required for Configuration and Installation

Before installation, collect the following information:

Required Information	Mandatory / Optional	Your Information / Notes
For Preinstallation Configuration		
Cisco ICFPP image location	Mandatory	

Required Information	Mandatory / Optional	Your Information / Notes
Cisco ICFPP OVA or QCOW2 image name	Mandatory	
VM name	Mandatory	
VMware datastore location	Mandatory for VMware	
Network / Port Profile for VM management	Mandatory	
KVM flavor name	Mandatory for OpenStack	
KVM Instance Security Group	Mandatory for OpenStack	
For Cisco ICFPP Installation		
Installation type: Standalone, Primary, or Service Node For OpenStack environments, you can install only in Standalone mode.	Mandatory	
Hostname	Mandatory	
Admin / root / ShellAdmin account password	Mandatory	
Static IP address For OpenStack environments, this must be a public IP address.	Mandatory	
Subnet mask	Mandatory	
Gateway IP address	Mandatory	
Primary node IP address	Mandatory only for service node installations.	
NFS server IP address	Optional ¹	
NFS directory to mount	Optional	
Domain name.	Optional	

Required Information	Mandatory/Optional	Your Information / Notes
DNS server IP address	Mandatory	
NTP server IP address or fully qualified domain name (FQDN)	Mandatory	

- ¹ If you do not configure NFS in a cluster deployment, template creation and VM migration can fail if a service node fails. If NFS is not available, you can configure a Cinder volume.



Installing Cisco ICFPP on VMware

- [Cisco ICFPP Software, page 13](#)
- [Workflow for VMware Environments, page 14](#)
- [Installing Cisco ICFPP on VMware, page 14](#)
- [Configuring the IP Address for Network Access, page 17](#)

Cisco ICFPP Software

The Cisco ICFPP software is available for download from [Cisco.com](https://www.cisco.com). For assistance, contact your Cisco representative.

The Cisco ICFPP software package (ICFP-dk9-2.3.1-pkg.zip), contains the following files:

File	Description
icfpp-2.3.1.ova	Cisco ICFPP OVA file. Use this file to install Cisco ICFPP in VMware environments. See Workflow for VMware Environments, on page 14 .
icfpp-2.3.1.qcow2	Cisco ICFPP QCOW2 file. Use this file to install Cisco ICFPP in OpenStack environments. See Workflow for OpenStack Environments, on page 19 .
README	README file. This file contains information about installing and using Cisco ICFPP.

The Cisco ICFPP software includes an evaluation license with support for 20 VMs. To view the license details in the GUI after you install Cisco ICFPP, choose **Administration > License**, and expand the entry in the **License Keys** table.

Workflow for VMware Environments

Cisco ICFPP should be implemented by all service providers that interface with Cisco Cisco Intercloud Fabric Director platforms. The only exceptions to this are Amazon EC2 and Windows Azure, which are available to Cisco Intercloud Fabric through their native public cloud APIs.

The high-level tasks involved in deploying Cisco ICFPP in a VMware environment are:

- 1 Confirm that you have met the installation requirements—See [Installation Requirements](#), on page 7.
- 2 Gather the required information—See [Information Required for Configuration and Installation](#), on page 9.
- 3 Install Cisco ICFPP—See [Installing Cisco ICFPP on VMware](#), on page 14.
- 4 If needed after the installation, configure the Cisco ICFPP IP address—See [Configuring the IP Address for Network Access](#), on page 17.
- 5 (Optional) Configure Cisco ICFPP virtual appliances for a multiple-node cluster—See [Configuring Cisco ICFPP for Clusters](#), on page 31.
- 6 Configure communications with Cisco Intercloud Fabric Director—See the *Cisco Intercloud Fabric Getting Started Guide*.

Installing Cisco ICFPP on VMware

This procedure describes how to install Cisco ICFPP in a VMware environment.



Note

We recommend that you configure additional storage for all Cisco ICFPP nodes. If additional storage is not configured, all VM images that are uploaded from Cisco Intercloud Fabric Director are stored on the node's local disk. If the node fails, one or both of the following can occur:

- Any images stored on the node are no longer available.
- If the node is part of a cluster, template creation and VM migration fail.

If NFS is not available, you can configure a Cinder volume as described in [Configuring a Cinder Volume](#), on page 35.

Before You Begin

- Set your keyboard to United States English.
- Unzip the Cisco ICFPP software package to obtain the OVA file and the README file.
- Copy the Cisco ICFPP OVA image to a location that is available from the VMware vSphere Client.
- Make sure that all requirements are met as specified in [System Requirements](#), on page 7.
- Collect the information required for the installation. See [Information Required for Configuration and Installation](#), on page 9.
- Review the README file for information related to Cisco ICFPP installation and operation.

Procedure

- Step 1** Using the VMware vSphere Client, log in to the vCenter server.
- Step 2** Choose the host on which to deploy the Cisco ICFPP virtual appliance.
- Step 3** Choose **File > Deploy OVF Template**.
- Step 4** In the wizard, provide the information as described in the following table:

Screen	Action
Source	Choose the Cisco ICFPP OVA using one of the following methods: <ul style="list-style-type: none"> • Browse to the location, choose the file, and click Open. • Deploy from a URL on your local area network. Replace <i>FQDN</i> with the IP address or the fully qualified domain name, and click Next.
OVF Template Details	Verify the details.
End User License Agreement	Read the agreement and click Accept .
Name and Location	<ol style="list-style-type: none"> 1 Enter a name for the virtual appliance. 2 Choose the VMware data center or host where Cisco ICFPP will reside.
Deployment Configuration	Choose the type of deployment: <ul style="list-style-type: none"> • Standalone—Used for single-node deployments. • Primary Node—Used for HA deployment in a multiple-node cluster. • Service Node—Used in cluster deployments for handling requests.
Storage	Choose the location in which to store the Cisco ICFPP files.
Disk Format	Choose the required format for the virtual appliance disks: <ul style="list-style-type: none"> • Thick Provision Lazy Zeroed—Allocates storage immediately in thick format. • Thick Provision Eager Zeroed—Allocates storage in thick format. Creating disks might take longer using this option. • Thin Provision—Allocates storage on demand as data is written to disk.
Network Mapping	Choose the required network.

Screen	Action
Properties	Address any errors that are indicated in red-colored text below a selection box.
Node Mode	Choose the type of deployment for this node: Standalone, Primary Node, or Service Node. The mode you choose should match the deployment type in the Deployment Configuration screen.
ICFPP Hostname	Enter the hostname for the Cisco ICFPP node.
ICFPP Password	Enter and confirm the password to use for the admin, root, and ShellAdmin account access.
Static IP Address	Enter the static IP address to use for the Cisco ICFPP node.
Static IP Subnet Mask	Enter the subnet mask to apply to the node IP address.
IP Gateway	Enter the gateway IP address.
Primary Node IP Address for Service Node	For service nodes only, enter the IP address of the primary node or the virtual IP address (VIP) of the HA pair for database access.
NFS Server IP Address	Enter the IP address for an NFS server. Note If you do not configure NFS in a multiple-node cluster deployment, template creation and VM migration can fail if a service node fails.
NFS Server Directory to Mount	NFS server directory to be mounted.
Domain Name	Enter the domain name for the node, such as cisco.com.
DNS Server IP Address	Enter the DNS server IP address.
NTP Server IP (FQDN or IP Address)	Enter the NTP server IP address or fully qualified domain name.
Ready to Complete	Review the deployment settings for accuracy.

Step 5 Click **Finish**. A progress indicator displays the task status until Cisco ICFPP is deployed. For additional information, right-click the VM in the VMware vSphere Client and choose **Open Console**.

Step 6 After Cisco ICFPP is successfully deployed, power on the virtual appliance.

What to Do Next

If needed, configure the Cisco ICFPP IP address for network address. For more information, see [Configuring the IP Address for Network Access](#), on page 17.

Configuring the IP Address for Network Access

After installing Cisco ICFPP in a VMware environment, you might need to configure the Cisco ICFPP IP address for network access.

The Cisco ICFPP IP address is configured during installation by using Open Virtualization Format (OVF) parameters. However, if the IP address is not configured correctly, you must configure the static IP address by using the ShellAdmin console options as described in this procedure.

Procedure

- Step 1** Using SSH, connect to the Cisco ICFPP ShellAdmin console by using the following information:
- Cisco ICFPP IP address
 - Username—shelladmin
 - Password—The password that you set when you installed Cisco ICFPP
- Step 2** At the ShellAdmin prompt, choose the **Configure Network Interface** option to configure the static IP address.
- Step 3** Enter **S** to configure a static IP address.
- Step 4** Enter the Ethernet interface that you want to configure, such as eth0 or eth1.
- Step 5** When prompted for the IP version, choose **IPv4**.
- Step 6** Enter the static IP address, netmask, and gateway IP address.
- Step 7** Enter **Y** to confirm the information.
The Cisco ICFPP virtual appliance reboots and displays a screen with the URL for accessing Cisco ICFPP.
- Step 8** (Optional) To verify that the change has been applied, log in to the ShellAdmin console and choose the **Display Network Details** option.
-



Installing Cisco ICFPP on OpenStack

- [Workflow for OpenStack Environments, page 19](#)
- [Installing Cisco ICFPP on OpenStack, page 19](#)
- [Configuring Cisco ICFPP for Cisco Intercloud Fabric Director, page 21](#)

Workflow for OpenStack Environments

Cisco ICFPP should be implemented by all service providers that interface with Cisco Cisco Intercloud Fabric Director platforms. The only exceptions to this are Amazon EC2 and Windows Azure, which are available to Cisco Intercloud Fabric through their native public cloud APIs.

The high-level tasks involved in installing and configuring Cisco ICFPP in an OpenStack environment are:

- 1 Confirm that you have met the installation requirements—See [Installation Requirements, on page 7](#).
- 2 Gather the required information—See [Information Required for Configuration and Installation, on page 9](#).
- 3 Configure OpenStack for Cisco ICFPP and launch a Cisco ICFPP instance—See [Installing Cisco ICFPP on OpenStack, on page 19](#).
- 4 Configure Cisco ICFPP for use with Cisco Intercloud Fabric Director—See [Configuring Cisco ICFPP for Cisco Intercloud Fabric Director, on page 21](#).
- 5 (Optional) Configure Cisco ICFPP virtual appliances for a multiple-node cluster—See [Configuring Cisco ICFPP for Clusters, on page 31](#).
- 6 Configure Cisco Intercloud Fabric Director for use with Cisco ICFPP—See the *Cisco Intercloud Fabric Getting Started Guide*.

Installing Cisco ICFPP on OpenStack

To install Cisco ICFPP on OpenStack, you must import an image, create a flavor, and launch an instance. This procedure describes how to complete these tasks.

The amount of time required for this procedure depends on the platform:

- If the platform does not support QCOW2, the procedure can take up to two hours to complete, depending on the amount of time it takes to upload the image and convert it from QCOW2 format to RAW.
- If the platform supports QCOW2, no conversion is required, and the procedure takes less time.

Before You Begin

- Download the Cisco ICFPP software package from cisco.com. For assistance, contact your Cisco representative.
- Unzip the downloaded file to obtain the QCOW2 file and the README file. For more information, see [Cisco ICFPP Software, on page 13](#).
- Review the README file for information related to installing and using Cisco ICFPP with OpenStack.
- Confirm that you have met the requirements in [System Requirements, on page 7](#).
- Gather the information identified in [Information Required for Configuration and Installation, on page 9](#).
- In OpenStack:
 - Confirm that you have admin privileges.
 - Create an OpenSource RC file (*name-openrc.sh*) in which you define your environmental variables and login credentials.
 - Create a project on which to install Cisco ICFPP.
 - Confirm that the Cinder service is up and running.
 - Configure a security group that allows traffic on ports 22, 80, 443, and 3306.

For more information about performing these operations in OpenStack, see docs.openstack.org.

Procedure

Step 1 In the shell from which you will enter **glance** commands, enter the following command:

```
source name-openrc.sh
```

Step 2 Copy the Cisco ICFPP image to the system running the **glance** CLI.

Step 3 Using the **glance** CLI, upload an image to the OpenStack server by entering the following command:

```
glance image-create --name icfpp-n.n.n --disk-format qcow2 --container-format bare --file ./icfpp-n.n.n.qcow2
```

where *icfpp-n.n.n* is the name of the Cisco ICFPP image, such as *icfpp-2.3.1*.

After the image has been uploaded, it appears in the OpenStack Dashboard Images table at **Admin > Images** or **Project > project > Manage Compute > Images & Snapshots**.

Step 4 In the OpenStack Dashboard, choose **Admin > Flavors**, and click **Create Flavor**.

Step 5 In the **Create Flavor** dialog box, enter the following information, and click **Create Flavor**:

- Name—Enter a flavor name.

- vCPUs—Enter **4**.
- RAM MB—Enter **8192**.
- Root Disk—Enter the desired disk size in gigabytes.
- Ephemeral Disk—Enter **0**.
- Swap Disk—Enter **0**.

Step 6 Choose **Project** > *project* > **Manage Compute** > **Volumes**, and click **Create Volume**.

Step 7 In the **Create Volume** dialog box, add a volume with the size 100 GB, and click **Create Volume**.

Step 8 In OpenStack, obtain the following information:

- Flavor ID
- Image ID
- Network ID

Step 9 At the command line, enter the following command to launch Cisco ICFPP:

```
nova boot --image image-id --flavor flavor-id  
--nic net-id=network-id --block-device-mapping vdb=volume-id  
icfpp-instance-name
```

A Cisco ICFPP instance is launched.

Configuring Cisco ICFPP for Cisco Intercloud Fabric Director

After you have installed Cisco ICFPP on an OpenStack server and launched a Cisco ICFPP instance, you can configure Cisco ICFPP for use with Cisco Intercloud Fabric Director.

Before You Begin

Confirm the following:

- Cisco ICFPP has been installed on an OpenStack server and an instance has been launched.
- You know the Cisco ICFPP public IP address.

Procedure

Step 1 In a browser, enter the public IP address assigned to the Cisco ICFPP instance and log in to the Cisco ICFPP GUI. The default credentials are:

- Username: admin
- Password: admin

Step 2 In the OpenStack dashboard, choose **Project** > *project* > **Access & Security**, and click the **API Access** tab.

Step 3 In the **API Endpoints** table, locate and make a note of the service endpoint Uniform Resource Identifier (URI) for the **Identity** service.

Step 4 In the Cisco ICFPP GUI, choose **Cloud Instances**, and click **Add**.

Step 5 In the **Add Cloud Instance** dialog box, provide the following information, and click **Add**:

Field	Description
Cloud Instance Name	The name of the cloud instance.
Type	The cloud instance type: Cisco or Custom.
Module Name	For a Cisco cloud instance type, choose the module name, such as OSP for OpenStack Platform. For a custom cloud instance type, enter the custom module name.
Image Conversion Support on Cloud	For OSP modules, indicate whether or not image conversion on the cloud is required.
First Boot Image Conversion Support	For OSP modules, indicate whether or not image conversion during VM boot on the cloud is required.
FTP Server Name	For Cisco Intercloud Services — V modules, the name of the FTP server.
Endpoint URI	The endpoint URI for the cloud instance.

Step 6 In the Cisco ICFPP GUI, choose **Tenants** > **All Tenants**, and click the **Accounts** tab.

Step 7 Click **Add**.

Step 8 In the **Add Tenant** dialog box, provide the following information, and click **Add**:

Field	Description
Tenant Name	Enter the tenant name. You cannot change the name after adding the tenant.
Cloud Instance Name	Choose the name of the cloud instance. You cannot change the cloud instance name after adding the tenant.
Enable Tenant Account	
Enabled	(Read-only) Indicates whether or not the tenant account is enabled. The account is enabled by default.
Org Name	For VMware vCloud Director clouds, enter the name of the organization to which the tenant belongs.

Field	Description
Resource Limits	
Max Servers	Enter the maximum number of servers provisioned for the tenant, including stopped VMs.
User Account	
Username	Enter the account username.
Email	Enter the account email address.
API Key	For CloudStack clouds, enter the API key for the tenant.
Secret Key	For CloudStack clouds, enter the Secret key for the tenant.

For information about configuring Cisco ICFPP for a multiple-node cluster, see [Configuring Cisco ICFPP for Clusters](#), on page 31.



CHAPTER 6

Upgrading Cisco ICFPP

- [Upgrading Standalone Nodes or Multiple-Node Clusters, page 25](#)
- [Supported Upgrade Paths, page 25](#)
- [Restarting Services Automatically, page 25](#)
- [Upgrading a Standalone Node, page 26](#)
- [Upgrading a Multiple-Node Cluster, page 27](#)

Upgrading Standalone Nodes or Multiple-Node Clusters

Cisco ICFPP enables you to upgrade standalone appliances and multiple-node clusters for bug fixes and updated adapters. For more information, see the following topics:

- [Supported Upgrade Paths, on page 25](#)
- [Restarting Services Automatically, on page 25](#)
- [Upgrading a Standalone Node, on page 26](#)
- [Upgrading a Multiple-Node Cluster, on page 27](#)

Supported Upgrade Paths

Cisco ICFPP 2.3.1 supports the following upgrade paths:

- OpenStack—Cisco ICFPP 2.2.1 to 2.3.1.
- VMware—Cisco ICFPP 2.2.1 or 2.2.1a to 2.3.1.

Restarting Services Automatically

Beginning with version 2.3.1, Cisco ICFPP includes a feature that automatically restarts Infra services when you upgrade Cisco ICFPP.

The first time that you upgrade Cisco ICFPP from 2.2.1 or 2.2.1a to 2.3.1 or higher, you must manually restart services. After you restart Infra services, the automatic service restart feature is enabled and you do not need to restart Infra services when you next upgrade Cisco ICFPP.

Upgrading a Standalone Node

This procedure enables you to apply Cisco bug fixes and upgrade adapters on a standalone node. To upgrade a multiple-node cluster, see [Upgrading a Multiple-Node Cluster](#), on page 27.

Before You Begin

- Obtain the Cisco ICFPP upgrade file (`icfpp-upgrade-2.3.1.tar.gz`) from [cisco.com](#). For assistance, contact your Cisco representative.
- Confirm that the upgrade file is accessible from the Cisco ICFPP virtual appliance.

Procedure

Step 1 In the Cisco ICFPP GUI, choose **Install > Adapters**, and click **Install**.

Step 2 In the **Install Adapter** dialog box, provide the information as described in the following table and then click **Upload**:

Field	Description
Adapter Type	Choose Cisco .
Adapter Name	This field displays CAPI by default. No input is required.
Adapter Description	Enter the desired description.
Adapter File	Browse to the Cisco ICFPP upgrade file and click Open .

Step 3 If you are upgrading from Cisco ICFPP 2.2.1 or 2.2.1a to Cisco ICFPP 2.3.1 or higher, complete the following steps:

- After the file has been uploaded, click **Submit**.
- Using SSH, log in to the ShellAdmin console for the virtual appliance.
- Choose the **Stop Services** option.
- Choose the **Start Services** option.

Step 4 If you are upgrading from Cisco ICFPP 2.3.1 to a higher version, a message is displayed stating that the upgrade will start in two minutes. After approximately two minutes, the upgrade is installed, the services automatically restart, and the GUI becomes unresponsive. To finish the upgrade, refresh the browser and log in to the Cisco ICFPP GUI.

Step 5 To verify that the upgrade was successful, click **About** in the GUI toolbar and confirm that the correct version is displayed.

The Product Version field displays the version using the format *version-build-patch* where:

- *version* is the product version, such as 2.3.1.

- *build* is the build number, such as 204.
- *patch* is the patch applied to the version and build, such as p208.

For example, you might see the version 2.3.1-204-p208.

Upgrading a Multiple-Node Cluster

Use this procedure to upgrade a multiple-node cluster for bug fixes and updated adapters. To upgrade a standalone Cisco ICFPP virtual appliance, see [Upgrading a Standalone Node, on page 26](#).

This procedure applies to multiple-node clusters with the following components and configuration:

- An HA pair consisting of two Cisco ICFPP virtual appliances that are configured as primary nodes.
- The HA pair is configured with one active node and one standby node.
- Additional Cisco ICFPP virtual appliances are configured as service nodes.

The workflow for upgrading a cluster includes the following high-level tasks:

- 1 Stop the virtual IP (VIP) service on the HA active node.
- 2 Monitor status while services fail over to the HA standby node.
- 3 Upgrade the current HA active node (originally the standby node).
- 4 Start the VIP service on the current HA standby node (originally the active node).
- 5 Stop the VIP service on the upgraded HA active node.
- 6 Monitor status while services fail over to the current HA standby node, making it the active node again.
- 7 Upgrade the current HA active node.
- 8 Start the VIP service on the current HA standby node.
- 9 Upgrade each service node.

The following procedure describes how to perform these tasks.

Before You Begin

- Obtain the Cisco ICFPP upgrade file (`icfpp-upgrade-2.3.1.tar.gz`) from Cisco.com. For assistance, contact your Cisco representative.
- Ensure that the upgrade file is accessible from the Cisco ICFPP virtual appliance.
- Confirm that HA has been configured on two Cisco ICFPP virtual appliances that are configured with the Primary Node role.

Procedure

- Step 1** Stop the VIP service on the HA active node as follows:
- Log in to the ShellAdmin console for the HA active node.
 - Choose **Setup HA**.
 - When asked if you want to reconfigure HA, enter **Y**.
 - Enter **C** to stop the VIP service.
 - Enter **Y** to confirm the action.
 - Press **Enter** to return to the ShellAdmin menu.
- Step 2** Log in to the ShellAdmin console for the HA standby node.
- Step 3** In the ShellAdmin console for the standby node, choose **Display Services Status** to monitor the following events as they occur:
- HA services fail over to the standby node in the HA pair.
 - Infra services start running on the standby node.
 - The GUI for the standby node becomes available for logging in.
- It can take a few minutes for the services to start and for the GUI of the standby node to be accessible from the browser.
- Note** The node that was originally the HA standby node becomes the HA active node.
- Step 4** Upgrade the currently active node of the HA pair as follows:
- Log in to the Cisco ICFPP GUI for the active node of the HA pair by using the management IP address of the node.
 - In the GUI, choose **Install > Adapters > Install**.
 - In the **Install Adapter** dialog box, provide the required information.
For more information about the fields in this dialog box, see [Upgrading a Standalone Node](#), on page 26.
 - Click **Upload**.
 - After the upload is complete, click **Submit**.
- Step 5** Do one of the following, depending on the Cisco ICFPP version:
- If you are upgrading from Cisco ICFPP 2.2.1 or 2.2.1a to 2.3.1, restart Infra services from the ShellAdmin console by first choosing **Stop Services** and then choosing **Start Services**.
 - If you are upgrading from Cisco ICFPP 2.3.1 to a higher version, the Infra services are restarted automatically and you can log in to Cisco ICFPP after approximately two minutes.
- Step 6** Verify that the HA active node was successfully upgraded as follows:
- Log in to the Cisco ICFPP GUI of the active node by using the management IP address of the node.
 - In GUI toolbar, click **About**.
 - Confirm that the correct version is displayed.
The version uses the format *version-build-patch*, such as 2.3.1-204-p208.
- Step 7** Restart the VIP service on the current HA standby node as follows:
- Log in to the ShellAdmin console for the current HA standby node.

- b) Choose **Setup HA**.
- c) When asked if you want to reconfigure HA, enter **Y**.
- d) Enter **D** to start the VIP service.
- e) Press **Enter** to return to the ShellAdmin menu.

Step 8 Stop the VIP service on the currently active node that was upgraded in Step 4 as follows:

- a) Log in to the Shell Admin console for the currently active node in the HA pair.
- b) Choose **Setup HA**.
- c) When asked if you want to reconfigure HA, enter **Y**.
- d) Enter **C** to stop the VIP service.
- e) Enter **Y** to confirm the action.
- f) Press **Enter** to return to the ShellAdmin menu.

Step 9 Log in to the ShellAdmin console for the standby node in the HA pair.

Step 10 In the ShellAdmin console for the standby node, choose **Display Services Status** to monitor the following events as they occur:

- HA services fail over to the standby node in the HA pair.
- Infra services start running on the standby node.
- The GUI for the standby node becomes available for logging in.

It can take a few minutes for the services to start and for the GUI of the standby node to be accessible from the browser.

Note The node that was previously the HA standby node becomes the HA active node.

Step 11 Upgrade the HA active node as follows:

- a) Using the management IP address instead of the virtual IP address for the HA pair, log in to the Cisco ICFPP GUI for the HA active node.
- b) Upgrade the node as described in Step 4.
- c) If needed, restart services as described in Step 5.
- d) Verify that the upgrade was successful as described in Step 6.

Step 12 Restart the VIP service on the HA standby node as follows:

- a) Log in to the ShellAdmin console for the HA standby node.
- b) Choose **Setup HA**.
- c) When asked if you want to reconfigure HA, enter **Y**.
- d) Enter **D** to start the VIP service.
- e) Press **Enter** to return to the ShellAdmin menu.

Step 13 Upgrade each service node in the cluster as follows:

- a) Log in to the Cisco ICFPP GUI for the service node.
- b) Upgrade the service node by uploading and submitting the upgrade package as described in Step 4.
- c) Do one of the following, depending on the Cisco ICFPP version:
 - If you are upgrading from Cisco ICFPP 2.2.1 or 2.2.1a to 2.3.1, restart Infra services from the ShellAdmin console by first choosing **Stop Services** and then choosing **Start Services**.
 - If you are upgrading from Cisco ICFPP 2.3.1 to a higher version, the Infra services are restarted automatically and you can log in to the service nodes after approximately two minutes.

Step 14 Verify that each service node upgraded successfully as follows:

- a) For each service node, refresh the browser and log in to the Cisco ICFPP GUI for the service node.
 - b) Click **About** in the GUI toolbar and confirm that the correct version is displayed.
The version uses the format *version-build-patch*, such as 2.3.1-204-p208.
-



Configuring Cisco ICFPP for Clusters

- [Workflow for Configuring Clusters, page 31](#)
- [Configuring a Primary Node, page 32](#)
- [Configuring a Service Node, page 33](#)
- [Configuring Additional Storage, page 34](#)
- [Configuring HA, page 37](#)
- [Configuring VIP Access for HA in OpenStack, page 39](#)
- [Moving from a Standalone Setup to a Cluster, page 42](#)
- [Restoring a Database onto an Existing HA Pair, page 43](#)
- [Monitoring HA Status, page 44](#)
- [Viewing HA Syslog Messages, page 45](#)

Workflow for Configuring Clusters

The following workflow describes the high-level tasks that are required to configure a multiple-node cluster.

Step	Task	Related Information
1.	Install a minimum of four Cisco ICFPP virtual appliances. The role that is assigned to each appliance during installation depends on whether you are using VMware or OpenStack.	Deployment Workflows, on page 5
2.	Configure two primary nodes.	Configuring a Primary Node, on page 32
3.	Configure two or more service nodes.	Configuring a Service Node, on page 33
4.	Configure additional storage.	Configuring Additional Storage, on page 34

Step	Task	Related Information
5.	Configure the two primary nodes for HA.	Configuring HA, on page 37
6.	(OpenStack only) Configure VIP access.	Configuring VIP Access for HA in OpenStack, on page 39
7.	Configure a load balancer for the service nodes in the cluster. Note The load balancer must be configured to persist sessions based on the PERSISTICFPP cookie that Cisco ICFPP issues.	Your load balancer documentation

Configuring a Primary Node

To configure a Cisco ICFPP virtual appliance that has been installed using the Standalone Mode role for a multiple-node cluster, you must first configure it as a primary node or service node by using the ShellAdmin console. This procedure describes how to configure a standalone node as a primary node. To configure a standalone node as a service node, see [Configuring a Service Node, on page 33](#).

Before You Begin

Install a Cisco ICFPP virtual appliance using the Standalone Mode role.

Procedure

-
- Step 1** Using SSH, log in to the ShellAdmin console of the standalone node that you want to configure as a primary node.
 - Step 2** At the ShellAdmin prompt, choose the **Change Node Role** option.
 - Step 3** When prompted, enter **Y** to change the node role.
 - Step 4** Enter **A** to configure the node as a primary node.
 - Step 5** Enter **Y** to confirm that you want to configure the node as a primary node. Information similar to the following is displayed:

```

user selected 'y'
Checking DB Status
  2399 ?      00:00:00 mysqld_safe
  2820 ?      00:04:21 mysqld
Configuring as Primary Node...
Stopping services before changing node role
Stopping the services...
Setting up current node as Primary node...
Enabling Remote Database access to ICFPP Service nodes
Checking the MySQL to be ready before enabling remote access to DB...
Waiting a maximum of 900 seconds for MySQL to be up on localhost

Trying a maximum of 900 seconds for enabling remote access to DB
Successfully enabled remote access for database

SUCCESS: Successfully changed node role to Primary Node

```

```

Stopping Database and restarting it for changes to take effect
Stopping database...
Database stopped...
Starting services that were previously stopped.
Starting the Database...
Starting the services...
In order for changes to take effect logout and log back in
Do you want to logout [y/n]?

```

- Step 6** Enter **Y** when prompted to log out.
You are logged out of the ShellAdmin console. When you log in again, the ShellAdmin menu will include options for configuring HA and viewing HA status.
-

Configuring a Service Node

To configure a Cisco ICFPP virtual appliance that has been installed using the Standalone Mode role for a multiple-node cluster, you must first configure it as a primary node or as a service node by using the ShellAdmin console. This procedure describes how to configure a standalone node as a service node. To configure a standalone node as a primary node, see [Configuring a Primary Node, on page 32](#).

Before You Begin

- Install a Cisco ICFPP virtual appliance using the Standalone Mode role.
- Obtain the IP address of a primary node in the cluster or the virtual IP address (VIP) of an HA pair in the cluster.
- Back up any data in the virtual appliance database that you want to keep. When the virtual appliance is reconfigured as a service node, the existing data will be deleted.

Procedure

- Step 1** Using SSH, log in to the ShellAdmin console of the standalone node that you want to configure as a service node.
- Step 2** At the ShellAdmin prompt, choose the **Change Node Role** option.
- Step 3** When prompted, enter **Y** to change the node role.
- Step 4** Enter **B** to configure the node as a service node.
- Step 5** Enter **Y** to confirm that you want to configure the node as a service node.
- Step 6** When asked if you want to continue, do one of the following:
- Enter **N** to stop the configuration so that you can back up the database.
 - Enter **Y** to confirm that you want to continue.

If you choose to continue, Cisco ICFPP confirms your choice.

- Step 7** When prompted, enter the IP address of the primary node or the VIP of the HA pair that the service node is to use.
Information similar to the following is displayed:

```
Configuring as Service Node...
```

```

Stopping services before changing node role
Stopping the services...
Setting up current node as ICFPP service node...with remote DB IP 123.45.1.60
Disabling Database service at startup

SUCCESS: Successfully changed node role to Service Node

Starting services that were previously stopped...
Starting the services...
In order for the changes to take effect, log out and log in again
Do you want to log out [y/n]?

```

- Step 8** Enter **Y** to log out.
The next time that you log in, the menu will include the options available for a service node.
-

Configuring Additional Storage

The default disk size of 100 GB for Cisco ICFPP is not sufficient for configuring Cisco ICFPP in a multiple-node cluster. As a result, you must add additional disk space before configuring a multiple-node cluster. You can use either NFS or a Cinder volume as described in the following topics:

- [Configuring NFS, on page 34](#)
- [Configuring a Cinder Volume, on page 35](#)

Configuring NFS

If you did not configure an NFS server for a Cisco ICFPP virtual appliance when you installed it, you can configure the appliance for NFS by using the ShellAdmin console.



Note

We recommend that you configure additional storage for all Cisco ICFPP nodes. If additional storage is not configured, all VM images that are uploaded from Cisco Intercloud Fabric Director are stored on the node's local disk. If the node fails, one or both of the following can occur:

- Any images stored on the node are no longer available.
- If the node is part of a cluster, template creation and VM migration fail.

If NFS is not available, you can configure a Cinder volume as described in [Configuring a Cinder Volume, on page 35](#).

Before You Begin

- Upload all images on the Cisco ICFPP virtual appliance to the cloud. If the images are not uploaded to the cloud, they are deleted when NFS is configured.
- Identify the NFS server IP address and the directory in which the files are to be stored.

Procedure

- Step 1** Using SSH, log in to the ShellAdmin console for the Cisco ICFPP virtual appliance that you want to configure for NFS.
- Step 2** Choose the **NFS Configuration** option.
Cisco ICFPP displays a menu with options for configuring, removing, and viewing an NFS configuration.
- Step 3** At the prompt, enter **A**.
Cisco ICFPP determines whether or not an NFS directory is mounted and displays the results:

```
Checking for mounted NFS directory...
NFS directory is not mounted
Note: Configuring NFS will delete any images that are not uploaded to the cloud! Proceed
[y/n]?
```

- Step 4** Enter **Y** to continue.
Cisco ICFPP determines whether or not an NFS IP address or NFS directory has been configured and then prompts you for input.
- Step 5** When prompted, enter the NFS server IP address and the NFS directory path.
Information similar to the following is displayed while NFS is configured:

```
Configuring NFS with : NFS Server IP=123.15.1.1, remote directory=/nfs/dir local mounting
point=/mnt/icfpp-images
Creating /mnt/icfpp-images directory.
Starting portmap and nfs services...
Starting portmap: [ OK ]
mount -t nfs 123.15.1.1:/icfpp-images /mnt/icfpp-images
May wait for mount up to 12-0 seconds..., please be patient...
Successfully mounted 123.15.1.1:/icfpp-images at /mnt/icfpp-images
Saving NFS Configuration
NFS IP address: 123.15.1.1
NFS Directory Path: /icfpp_images
Saved NFS Configuration
Setting up images directory to use NFS
Image directory setup to NFS done
Press Return to continue
```

- Step 6** Press **Enter** to return to the ShellAdmin menu.
To view or remove the NFS configuration, choose the **NFS Configuration** option in the ShellAdmin menu, and then choose the appropriate option from the NFS menu.

Configuring a Cinder Volume

The default disk size of 100 GB for the Cisco ICFPP virtual appliance is not sufficient for configuring Cisco ICFPP in a multiple-node cluster. If you do not have access to an NFS server, you can increase the disk size by creating additional Cinder volumes. Cinder volumes that you create are formatted as physical disks and then combined to form a logical volume that can be mounted on the VM in a specific location.

Before You Begin

- Configure a Cisco ICFPP virtual appliance as a service node by using the ShellAdmin console. For more information, see [Configuring a Service Node](#), on page 33.

- If you have not already done so, configure the root user password for the Cisco ICFPP service node. For more information, see the "Using Cisco ICFPP ShellAdmin Commands" chapter in the [Cisco Intercloud Fabric Provider Platform Administrator Guide](#).
- Collect the following information:
 - Cloud credentials—The username and password for the project in OpenStack.
 - Cloud URL—Obtain the cloud URL as follows:
 - 1 In the OpenStack dashboard, choose **Project** > *project* > **Access & Security**, and click the **API Access** tab.
 - 2 In the **API Endpoints** table, locate the **Identity** service and note the service endpoint URL.
 - Cisco ICFPP instance ID—Obtain the Cisco ICFPP instance ID as follows:
 - 1 In the OpenStack dashboard, choose **Project** > *project* > **Instances**.
 - 2 In the list of instances, locate Cisco ICFPP and click the hyperlinked instance name. The **Instance Detail** page is displayed.
 - 3 In the Overview tab, locate and note the instance ID.

Procedure

- Step 1** Using SSH, log in to the ShellAdmin console of the Cisco ICFPP service node.
- Step 2** At the ShellAdmin prompt, choose **Cinder Storage Configuration**.
- Step 3** When prompted, enter **Y** and enter the root password.
- Step 4** At the Cinder Storage Configuration menu prompt, choose **Deploy Fresh Storage**. Cisco ICFPP prompts you for information so that it can configure the storage.
- Step 5** Enter the following information:
- Cloud username and password
 - OpenStack project name
 - Cloud URL
 - Cisco ICFPP instance ID
 - Required storage size in GB
 - Required volume size in GB

Note Cinder storage configuration supports a volume with a maximum of 2 TB for each service node.

Information similar to the following is displayed while Cisco ICFPP creates and formats the volume. You do not need to restart the Cisco ICFPP virtual appliance.

```
Cloud user name:- abc1-de2.gen
Enter password:
Project Name:- ABC-DEV1
Cloud URL: [e.g. https://us-texas-3.cloud.abc.com:5000/v2.01]:-
```

```

https://us-texas-3.cloud.abc.com:5000/v2.0
ICFPP Instance ID:- 75c8c226-b22c-4041-ab5c-7e7fd544c3b
Expected storage size[GB]:- 10
Expected volume size[GB]:- 10
Deploying fresh storage

****Creating volumes****

****Attaching volumes****

****Formatting volumes and creating logical volumes****

****Validating final state****
true
Executed successfully!

```

Step 6 If needed, you can do either of the following from the Cinder Storage Configuration menu:

- To configure additional storage, choose **Add additional storage to existing storage**.
- To delete storage, choose **Cleanup deployed storage**.

Configuring HA

After you deploy Cisco ICFPP virtual appliances, you can configure them for high availability (HA) by using the ShellAdmin console.

When configuring HA:

- Configure the active node and standby node concurrently as described in this procedure.
- The database on the standby node is deleted when the HA pair is configured.

Before You Begin

- Deploy or configure two Cisco ICFPP virtual appliances as primary nodes:
 - To deploy a Cisco ICFPP virtual appliance with the Primary Mode role, see [Deployment Workflows](#), on page 5.
 - To configure an existing Cisco ICFPP virtual appliance as a primary node, see [Configuring a Primary Node](#), on page 32.
- Identify a virtual IP (VIP) address for the HA pair.
- Determine which node will be the active node and which node will be the standby node.
- On the node that will be the standby node, move any existing data that you want to save to another location.

Procedure

- Step 1** Using SSH, log in to the ShellAdmin console of the node that will be the active node for the HA pair.
- Step 2** At the ShellAdmin prompt, choose the **Setup HA** option and press **Enter**.

A warning is displayed stating that the contents of the database on the standby node will be deleted.

Step 3 When prompted, enter **Y** to configure the node for HA.

Step 4 Enter **A** to configure the node as the active node.

Step 5 When prompted, enter **Y** to configure the node as the active node.
Cisco ICFPP detects and displays the IP address of the current node.

Step 6 Enter **Y** to confirm the node IP address.

Step 7 Enter the standby node IP address.

Step 8 Enter the VIP to use for the IP pair.
Information similar to the following is displayed:

```
-----
HA Configuration Information:
-----
This node will be configured as active node
Active Node IP address: 123.45.1.61
Standby Node IP address: 123.45.1.62
Virtual IP address:      123.45.1.60
-----
Proceed with setting up HA with above configuration [y/n]:
```

Step 9 Enter **Y** to confirm the configuration and proceed or **N** to change the values. If you choose to proceed, Cisco ICFPP displays progress messages while it configures the active node for HA.

Step 10 While Cisco ICFPP is configuring the active node for HA, log in to the ShellAdmin console of the node that will be the standby node for the HA pair.

Step 11 At the ShellAdmin prompt, choose the **Setup HA** option and press **Enter**.

Step 12 Enter **Y** to configure the node for HA.

Step 13 Enter **B** to configure the node as the standby node.

Step 14 When prompted, enter **Y** to configure the node as the standby node.
Cisco ICFPP detects and displays the IP address of the current node.

Step 15 Enter **Y** to confirm the node IP address.

Step 16 Enter the active node IP address.

Step 17 Enter the VIP to use for the HA pair.
Information similar to the following is displayed:

```
-----
HA Configuration Information:
-----
This node will be configured as standby node
Active Node IP address: 123.45.1.61
Standby Node IP address: 123.45.1.62
Virtual IP address:      123.45.1.60
-----
Proceed with setting up HA with above configuration [y/n]:
```

Step 18 Enter **Y** to confirm the configuration.
Cisco ICFPP displays progress messages while it configures the standby node for HA and synchronizes the database information on both nodes.

Step 19 When prompted, press **Enter** to return to the ShellAdmin menu.

What to Do Next

For OpenStack environments, continue with [Configuring VIP Access for HA in OpenStack](#), on page 39.

Configuring VIP Access for HA in OpenStack

After Cisco ICFPP primary nodes are configured for HA, the virtual IP address (VIP) is used in the event of failover. However, OpenStack Neutron does not allow a host to accept packets with an IP address in the packet header that does not match the destination host IP address. As a result, packets sent to the VIP do not reach the node to which the VIP is assigned. To allow the packets to reach HA pair, the VIP must be added as an allowed address for both nodes (active and standby) in the HA pair.

This procedure describes how to configure VIP access on the nodes in the HA pair by using the OpenStack **neutron port-update** command. For more information, see the OpenStack documentation at docs.openstack.org.

Before You Begin

- Confirm that HA has been configured on two Cisco ICFPP primary nodes in an OpenStack environment.
- Confirm that you have access to the OpenStack Neutron command line tool.

Procedure

Step 1 Obtain a list of networks by entering the following command:

```
$ neutron net-list
```

Information similar to the following is displayed:

id	name	subnets
2d84eaa4-8b81-4dc8-9897-dd8ef4719f8b	public-direct-600	10.203.28.0/23
3e0b77fe-fc66-4913-bc58-7f62d4ab247a		
5c2f73a9-4e2f-498c-8244-6aefe5129fdd		10.203.50.0/23
ba29165f-c88a-496a-9adc-99ee90407ebe		10.203.24.0/23
d5b69780-aefb-42a6-8ba5-aaf405fb36a0		10.203.30.0/24
b5d8d461-74d7-45a4-alf0-f7ac96586bd5	Net1	
c0921b42-2896-4b32-b33e-f54db9e5a3d6		192.168.0.0/24
ca80ff29-4f29-49a5-aa22-549f31b09268	public-floating-601	
0cfde3f1-e28b-4b87-8095-e0014b0ee573		
348a808d-ce64-43bc-a9d9-c20e52d2ac06		
3784170e-5d7f-48b4-b63d-aab4a0fef769		
ff95095f-89f0-4005-b709-70a75212d73c	icfpp-ha-123-network	
1099b814-05d9-4da0-93d1-06167db4891f		192.168.1.0/24

Step 2 Obtain a list of ports on the network on which the active and standby nodes in the HA pair are deployed by entering the following command:

```
$ neutron port-list -- --network_id=net_id
```

where *net_id* is the identifier for the required network. In this example, the network name is *icfpp-ha-123-network*.

```
$ neutron port-list -- --network_id=ff95095f-89f0-4005-b709-70a75212d73c
```

Information similar to the following is displayed:

id	name	mac_address	fixed_ips
4a439cf1-b95e-49ba-a8d6-0b03a8142dd2		fa:16:3e:f6:f8:a9	{"subnet_id": "1099b814-05d9-4da0-93d1-06167db4891f", "ip_address": "192.168.1.12"}
93d0a69a-7bb8-4719-9ed7-63c10accd78b		fa:16:3e:1f:7f:d2	{"subnet_id": "1099b814-05d9-4da0-93d1-06167db4891f", "ip_address": "192.168.1.11"}
9d626a64-ee7c-410b-ae00-661dd275de79		fa:16:3e:61:81:4b	{"subnet_id": "1099b814-05d9-4da0-93d1-06167db4891f", "ip_address": "192.168.1.14"}
cf56fd7b-2896-4e06-b520-1d2258ad6158		fa:16:3e:ab:27:ca	{"subnet_id": "1099b814-05d9-4da0-93d1-06167db4891f", "ip_address": "192.168.1.13"}
d7457d29-44ba-46ef-b47a-4b94c9199902		fa:16:3e:ad:d0:e9	{"subnet_id": "1099b814-05d9-4da0-93d1-06167db4891f", "ip_address": "192.168.1.15"}

Step 3 In the output of the previous step, locate the port ID for the active node.

Step 4 Update the port so that it accepts traffic from the VIP by entering the following command:

```
$ neutron port-update active-port-id --allowed_address_pairs list=true type=dict ip_address=vip
```

where:

- *active-port-id* is the port ID of the active node.
- *vip* is the virtual IP address for the HA pair.

For example, if the IP address of the active node is 192.168.1.11 and the VIP is 192.168.1.10, the command resembles the following:

```
$ neutron port-update 93d0a69a-7bb8-4719-9ed7-63c10accd78b --allowed_address_pairs list=true type=dict ip_address=192.168.1.10
```

Step 5 View the port details and confirm that the *allowed_address_pairs* field lists the VIP by entering the following command:

```
$ neutron port-show active-port-id
```

where *active-port-id* is the identifier for the port configured in the previous step.

Using the current example, the command and results resemble the following:

```
$ neutron port-show 93d0a69a-7bb8-4719-9ed7-63c10accd78b
```

Field	Value
admin_state_up	True
allowed_address_pairs	{"ip_address": "192.168.1.10", "mac_address": "fa:16:3e:1f:7f:d2"}

```

| device_id          | b7b8eeb5-70ad-49ac-a3b4-6d8a144293a2
| device_owner      | compute:alln01-1-csi
| extra_dhcp_opts   |
| fixed_ips         | {"subnet_id": "1099b814-05d9-4da0-93d1-06167db4891f", "ip_address":
"192.168.1.11"}
| id                | 93d0a69a-7bb8-4719-9ed7-63c10accd78b
| mac_address       | fa:16:3e:1f:7f:d2
| name              |
| network_id        | ff95095f-89f0-4005-b709-70a75212d73c
| security_groups   | f995d22f-edb8-47c0-9aff-6339a15fb5be
| status            | ACTIVE
| tenant_id         | b1436740f8db42e39904ee9779f67eb8
+-----+

```

Step 6 Configure the standby node to accept VIP traffic by entering the following command:

```
$ neutron port-update standby-port-id --allowed_address_pairs list=true type=dict ip_address=vip
```

where:

- *standby-port-id* is the port ID of the standby node.
- *vip* is the virtual IP address for the HA pair.

Step 7 View the port details for the standby node and confirm that the `allowed_address_pairs` field lists the VIP:

```
$ neutron port-show standby-port-id
```

Step 8 (Optional) Complete the following steps to configure the VIP so that it is accessible from an external network and so that the VIP uses a floating IP address:

a) Configure a port corresponding to the VIP by entering the following command:

```
$ neutron port-create --fixed-ip ip_address=ip --security-group security-group network-name
```

where:

- *ip* is the fixed IP address for the port.
- *security-group* is the name of the security group to use for this port.
- *network-name* is the name of the network to which the port belongs.

Using the current example, the command and results resemble the following:

```
$ neutron port-create --fixed-ip ip_address=192.168.1.10 --security-group default
icfpp-ha-123-network
```

Created a new port:

Field	Value
admin_state_up	True
allowed_address_pairs	
device_id	
device_owner	
fixed_ips	{"subnet_id": "1099b814-05d9-4da0-93d1-06167db4891f", "ip_address": "192.168.1.10"}
id	ea35e2a9-1b45-4b05-b345-f4758e490052
mac_address	fa:16:3e:df:e9:69
name	
network_id	ff95095f-89f0-4005-b709-70a75212d73c
security_groups	f995d22f-edb8-47c0-9aff-6339a15fb5be
status	DOWN
tenant_id	b1436740f8db42e39904ee9779f67eb8

- b) In the OpenStack Horizon GUI, associate a floating IP address with the port to which the fixed IP address is assigned.

Moving from a Standalone Setup to a Cluster

Cisco ICFPP enables you to move from a standalone configuration to a cluster. Moving from a standalone configuration to a cluster involves moving the database contents from the existing standalone node to the active HA node in the cluster as described in this procedure.

After moving the database contents, you can configure and test the cluster setup without modifying or affecting the standalone setup. For more information about configuring a multiple-node cluster, see [Deployment Workflows](#), on page 5.

Before You Begin

- Obtain the FTP server IP address and login credentials for backing up and restoring the database.
- Confirm that the standalone node and both nodes in the HA pair are running the same version of Cisco ICFPP.

Procedure

- Step 1** In the ShellAdmin console for the standalone node, back up the existing database as follows:
- Choose **Stop Services** to stop the Infrastructure Manager services.
 - Choose **Backup Database**.
 - Choose **Start Services**.
- Step 2** Deploy or configure two primary nodes by using any of the following methods:
- For VMware environments, deploy two new Cisco ICFPP virtual appliances using the Primary Node role. For more information, see [Installing Cisco ICFPP on VMware, on page 14](#).
 - For OpenStack environments, deploy two new Cisco ICFPP virtual appliances using the Standalone Node role and then configure the appliances as primary nodes. For more information, see [Installing Cisco ICFPP on OpenStack, on page 19](#).
 - Configure existing Cisco ICFPP virtual appliances using the Standalone Node role as primary nodes. For more information, see [Configuring a Primary Node, on page 32](#).
- Step 3** Restore the backed-up database from Step 1 onto one of the primary nodes:
- In the primary node ShellAdmin console, choose **Stop Services** to stop the Infrastructure Manager services.
 - Choose **Restore Database**.
 - Choose **Start Services**.
- Step 4** In the ShellAdmin console, configure the two primary nodes as an HA pair.
- Note** You must configure the primary node on which the database was restored as the active node in the HA pair. If you configure it as the standby node, the database on that node is deleted. For more information, see [Configuring HA, on page 37](#).
- Step 5** Configure service nodes for the cluster. For more information, see [Configuring a Service Node, on page 33](#).
-

Restoring a Database onto an Existing HA Pair

Cisco ICFPP enables you to configure an HA pair and then restore a database from an existing standalone node to the HA pair.



Note You must stop and start services in the sequence described in this procedure to successfully restore the database on the HA pair.

Before You Begin

- Confirm that the standalone node and both nodes in the HA pair are running the same version of Cisco ICFPP.
- Back up the required database from a standalone node onto an FTP server.
- Identify the active node in the HA pair on which to restore the backed-up database.

Procedure

- Step 1** Stop the VIP service on the current standby node in the HA pair as follows:
- Log in to the ShellAdmin console for the current standby node.
 - Choose the **Setup HA** option.
 - When asked if you want to reconfigure HA, enter **Y**.
 - Enter **C** to stop the VIP service.
 - Enter **Y** to confirm the action.
 - Press **Enter** to return to the ShellAdmin menu.
- Step 2** Stop the VIP service on the current active node in the HA pair as follows:
- Log in to the ShellAdmin console for the current active node.
 - Choose the **Setup HA** option.
 - When asked if you want to reconfigure HA, enter **Y**.
 - Enter **C** to stop the VIP service.
 - Enter **Y** to confirm the action.
 - Press **Enter** to return to the ShellAdmin menu.
- Stopping the VIP service on the active node in an HA pair automatically stops the Infrastructure Manager services if they are running.
- Step 3** On the active node in the HA pair, restore the database backup obtained from the standalone node as follows:
- In the ShellAdmin console for the active node, choose **Restore Database**.
 - When prompted, enter the FTP server IP address and login credentials.
 - Enter the path and filename for the backed up database file on the FTP server.
 - Follow the onscreen prompts to complete the process.
- Step 4** Restart the VIP service on the active node as follows:
- In the ShellAdmin console for the active node, choose **Setup HA**.
 - When asked if you want to reconfigure HA, enter **Y**.
 - Enter **D** to start the VIP service.
 - Press **Enter** to return to the ShellAdmin menu.
- Starting the VIP service on the active node in an HA pair automatically starts the Infrastructure Manager services on that node.
- Step 5** Restart the VIP service on the standby node in the HA pair as follows:
- In the ShellAdmin console for the standby node, choose **Setup HA**.
 - When asked if you want to reconfigure HA, enter **Y**.
 - Enter **D** to start the VIP service.
 - Press **Enter** to return to the ShellAdmin menu.
-

Monitoring HA Status

After configuring Cisco ICFPP for HA, you can view the configuration details, check the status of the active and standby nodes, and view detailed replication status.

Procedure

Step 1 Log in to the ShellAdmin console for one of the nodes in the HA pair.

Step 2 At the menu prompt, choose **Display HA Status**.
Information similar to the following is displayed:

```
Configured HA role for this node is: Active
Current HA role for this node is: Active
HA Configuration properties for this node are:
ACTIVE_IP_ADDRESS=123.16.1.30
STANDBY_IP_ADDRESS=123.16.1.3
VIRTUAL_IP_ADDRESS=123.16.1.25

IP address of this node is: 123.16.1.30
Checking if Virtual IP Address is reachable...OK
Virtual IP Address service status on this node...OK
Checking DB replication from 123.16.1.30 to 123.16.1.3...OK
Checking DB replication from 123.16.1.3 to 123.16.1.30...OK
```

Do you want to view detailed replication status ? [y/n]

Step 3 To view detailed information, enter **Y** and press **Enter**.
Information similar to the following is displayed:

```
Slave_IO_State : Waiting for master to send event
Master_Host : 123.16.1.3
Master_User : replicator
Master_Port : 3306
Connect_Retry : 60
Master_Log_File : mysql-bin.000002
Read_Master_Log_Pos : 645644
Relay_Log_File : mysqld-relay-bin.000004
Relay_Log_Pos : 361
Relay_Master_Log_File : mysql-bin.000002
Slave_IO_Running : Yes
Slave_SQL_Running : Yes
Replicate_Do_DB :
Replicate_Ignore_DB :
```

...

Step 4 Use your arrow keys to scroll through the information, and enter **Q** to return to the menu.

Viewing HA Syslog Messages

After configuring Cisco ICFPP for HA, Cisco ICFPP checks HA status every five minutes. Any warning or failure messages that are issued are included in the log file for syslog messages. This log file commonly resides in `/var/log/` with the name `messages`. To view these messages, log in as root and use a text editor as described in this procedure.

Procedure

- Step 1** In the ShellAdmin console, choose the **Log in as Root** option.
- Step 2** Enter **Y** to confirm the login request, and enter the root account password at the prompt.
- Step 3** Enter the following command to view the contents of the log file:

```
vi /directory-path/filename
```

where *directory-path* is location of the log file and *filename* is the name of the log file. For example, you might enter the following:

```
vi /var/log/messages
```

- Step 4** To identify messages that pertain to HA, look for entries that contain the string `ICFPP HA` as shown in the following example:

```
Mar 13 03:27:13 localhost logger: ICFPP HA: MySQL replication from 123.45.67.8 to 123.45.67.9
is in WARN state
Mar 13 03:27:13 localhost logger: ICFPP HA: Please use shelladmin to check HA status details
Mar 13 03:27:13 localhost logger: ICFPP HA: MySQL replication from 122.33.44.5 to 122.33.44.6
is in WARN state
Mar 13 03:27:13 localhost logger: ICFPP HA: Please use shelladmin to check HA status details
```

- Step 5** Address any HA-related messages as needed.
-



Configuring VMware vCloud Director for Cisco ICFPP

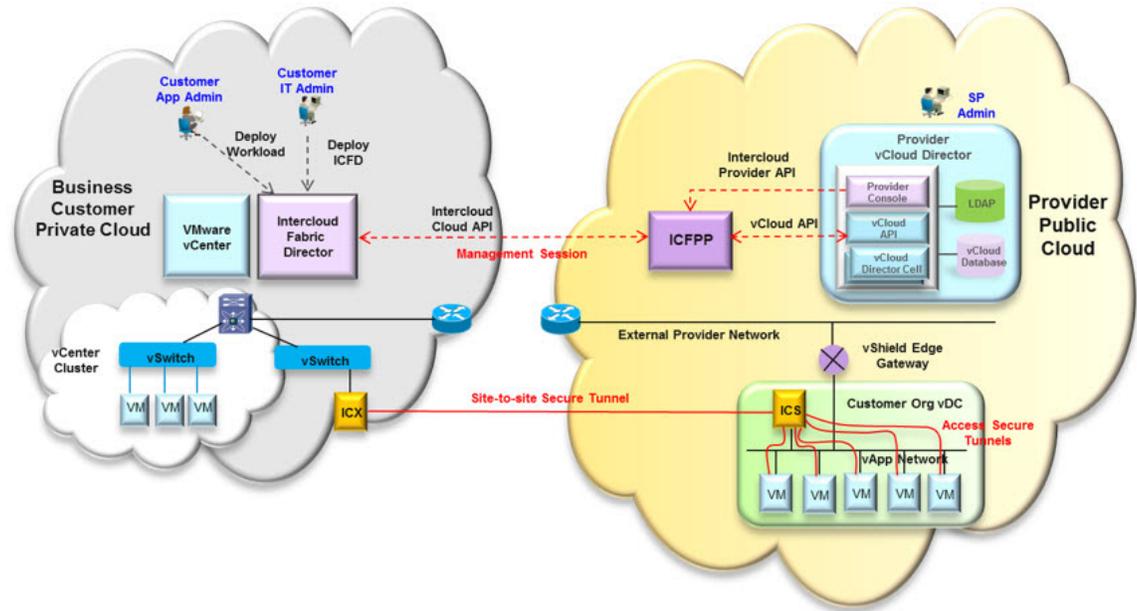
- [Configuring VMware vCloud Director, page 47](#)
- [Workflow for Integrating VCD with Cisco ICFPP, page 50](#)
- [Creating an External Network, page 51](#)
- [Adding a vShield Edge Gateway on an Org VDC, page 52](#)
- [Creating an Org VDC Internal Network, page 53](#)
- [Creating a Catalog, page 55](#)
- [Verifying NAT and Firewall Service Configuration, page 56](#)

Configuring VMware vCloud Director

Installing Cisco ICFPP at a cloud provider site enables you to support a hybrid cloud environment with Cisco Intercloud Fabric for Business. For VMware vCloud Director (VCD) environments, Cisco ICFPP includes a built-in VCD adapter that enables Cisco ICFPP to integrate with the VCD platform. This VCD-Cisco ICFPP integration can be viewed as the infrastructure that binds the enterprise virtualization platform, such as VMware vCenter, to the provider cloud platform, VCD.

The following illustration depicts how Cisco Intercloud Fabric Director interfaces with the provider VCD platform through Cisco ICFPP.

Figure 2: VCD and Cisco Intercloud Fabric Integration Architecture



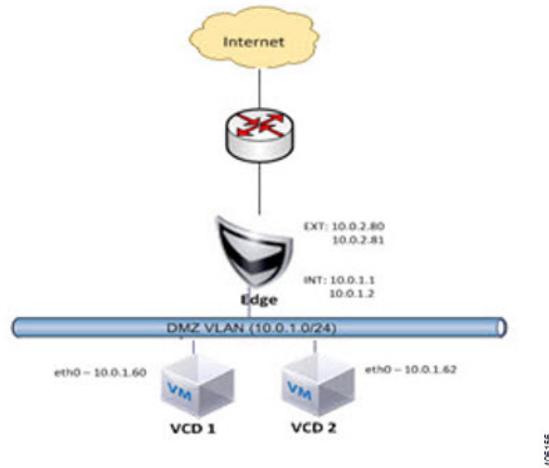
The secure site-to-site tunnel illustrated in the image is created between an Intercloud Fabric Switch (ICS) on the provider cloud and an Intercloud Fabric Extender (ICX) on the private cloud. In addition to providing secure communications between the private and provider clouds, this site-to-site tunnel enables Cisco Intercloud Fabric Secure Extender to integrate with VCD for each tenant network.

Before the ICS and ICX can communicate via the Internet, you must:

- Assign a public IP address to the ICS so that the ICX can reach the ICS.
- Ensure that the vShield Edge Gateway provides NAT functionality so that the ICS can connect to the Internet.

The following figure shows an example deployment:

Figure 3: vShield Edge Gateway Deployment Example

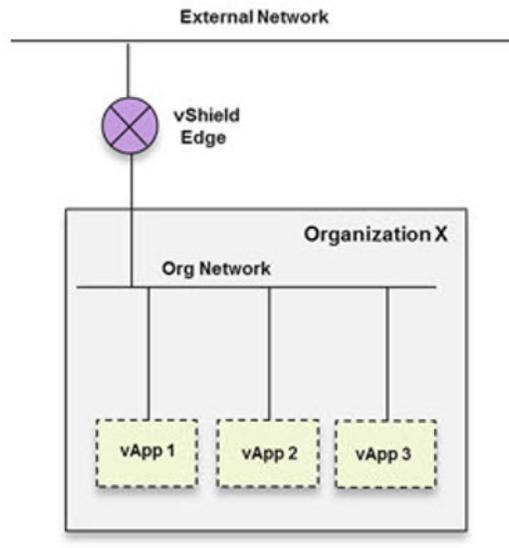


A vShield Edge Gateway is an interconnecting appliance which provides many edge network service features, including:

- NAT
- Firewall
- Load-balancer
- IPsec VPN
- DHCP
- Static route

The following figure shows how Organization X connects the Org Network to an external network through a vShield Edge Gateway and directly to vApp networks.

Figure 4: VCD Networking Model



426157

Workflow for Integrating VCD with Cisco ICFPP

To integrate VCD with Cisco ICFPP, you must provision certain infrastructure resources in the target VCD platform. The following table identifies the tasks required to provision these resources:

Step	Task	Related Information
1.	Ensure that the following prerequisites are met: <ul style="list-style-type: none"> • VCD version 5.5 is installed. • You have access to the VCD system administrator account. 	VMware VCD documentation
2.	Create an external network.	Creating an External Network, on page 51
3.	Deploy the vShield Edge Gateway.	Adding a vShield Edge Gateway on an Org VDC, on page 52
4.	Create an Org VDC network.	Creating an Org VDC Internal Network, on page 53
5.	Create a catalog.	Creating a Catalog, on page 55
6.	Ensure that NAT and firewall services are configured on the vShield Edge Gateway.	Verifying NAT and Firewall Service Configuration, on page 56

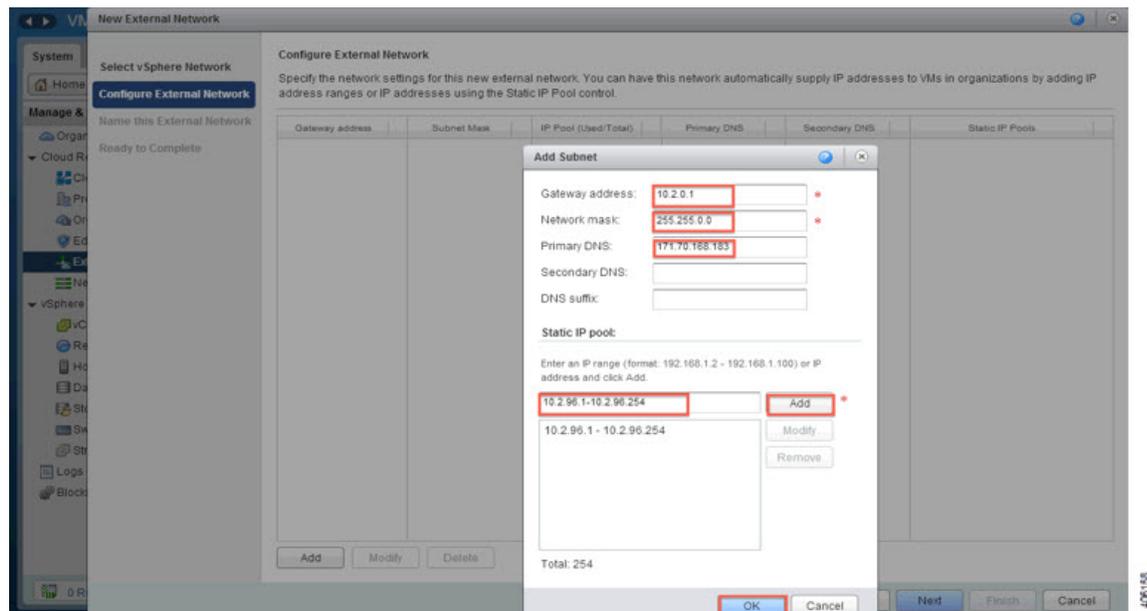
For additional information on any of these topics, see your VMware documentation.

Creating an External Network

This procedure describes how to create an external network in a virtual data center (VDC).

Procedure

- Step 1** Log in to the VCD GUI as system administrator.
- Step 2** Choose **System > Manage & Monitor > Cloud Resources > External Networks**.
- Step 3** In the **External Networks** pane, click **Add**.
The **New External Network** wizard opens, guiding you through the configuration process.
- Step 4** In the **Select vSphere Network** screen, choose the VDC vCenter and the DVS port group created for the vSphere management network, and click **Next**.
- Step 5** In the **Configure External Network** screen, click **Add**.
- Step 6** In the **Add Subnet** dialog box, enter the following information for the external network:
 - Gateway IP address
 - Network mask
 - DNS server IP address
 - Static IP address or IP address range



- Step 7** In the **Name this External Network** screen, enter a name for the external network, and click **Next**.
- Step 8** In the **Ready to Complete** screen, review the content for accuracy and click **Finish**.
The newly created external network is displayed in the **External Networks** pane.
-

Adding a vShield Edge Gateway on an Org VDC

You must add a vShield Edge Gateway to integrate the Provider VDC and Org VDC with Cisco ICFPP.

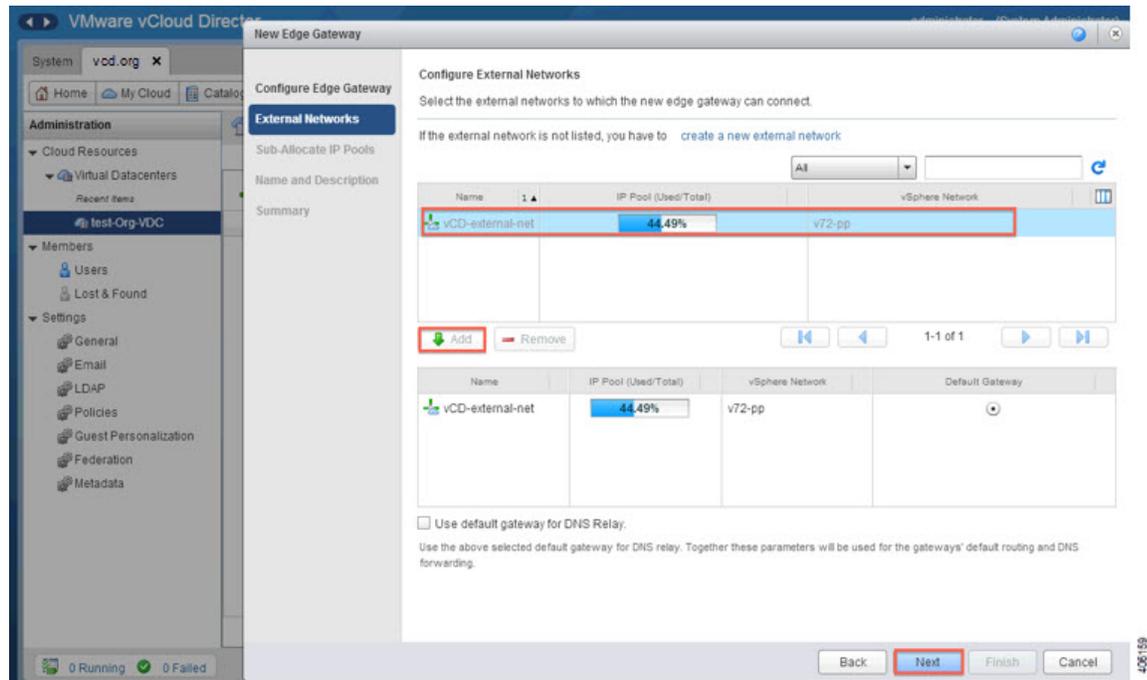
Before You Begin

Confirm that the following have been configured:

- A Provider VDC
- An Org VDC
- An external network

Procedure

- Step 1** In the VCD GUI, choose **System > Manage & Monitor > Cloud Resources > Organization VDCs**.
- Step 2** In the **Organization VDCs** table, double-click the Org VDC where the vShield Edge Gateway is to be added. The screen is refreshed with information about the selected VDC.
- Step 3** Choose the **Edge Gateways** tab and click **Add**.
The **New Edge Gateway** wizard opens, guiding you through the configuration process.
- Step 4** In the **Configure Edge Gateway** screen, configure the vShield Edge Gateway for connectivity with the external network as follows, and then click **Next**:
- Choose the required edge gateway configuration: Compact, Full, or Full-4.
 - If the edge gateway is to be configured for HA, check the **Enable High Availability** check box.
 - In the **Advanced Options** section, check the **Sub-Allocate IP Pools** check box.
- Step 5** In the **External Networks** screen, choose the external network that you created in [Creating an External Network, on page 51](#) and click **Add**. If the external network is not listed, create a new external network.



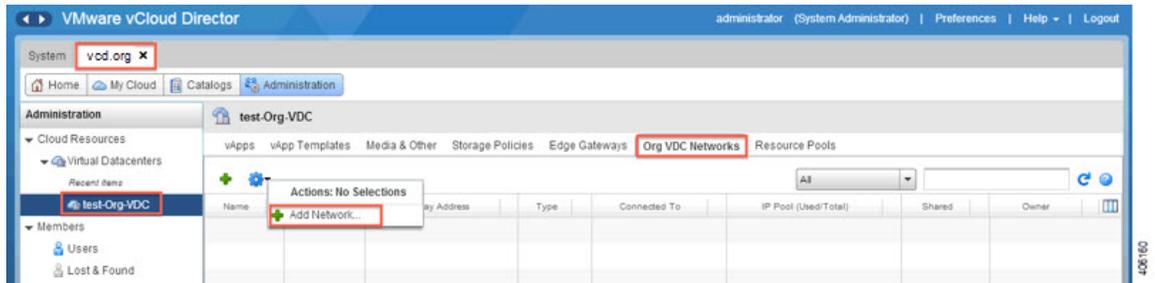
- Step 6** After the external network is added to the list of networks in the lower portion of the screen, click **Next**.
- Step 7** In the **Sub-Allocate IP Pools** screen, identify the range of IP addresses allocated for each externally-connected interface on the external network, and click **Next**.
- Step 8** In the **Name and Description** screen, enter the edge gateway name and description, and then click **Next**.
- Step 9** In the **Summary** screen, review the information for accuracy and click **Finish**.

Creating an Org VDC Internal Network

Use this procedure to create an internal network for the Org VDC.

Procedure

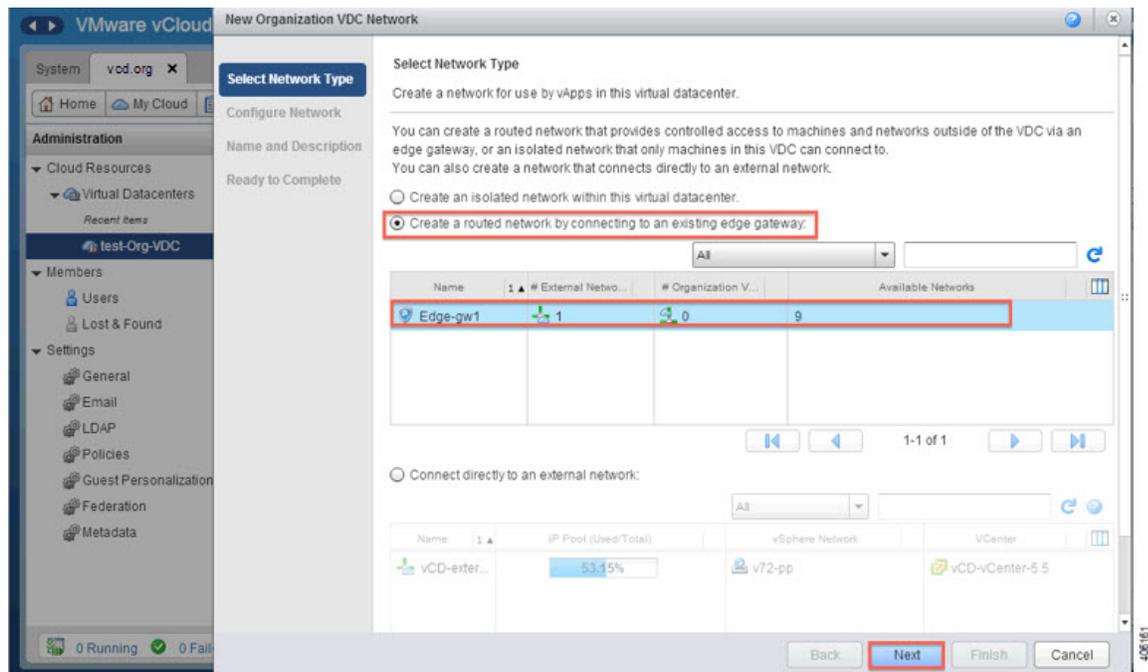
- Step 1** In the VCD GUI, choose **System > Manage & Monitor > Cloud Resources > Organization VDCs**.
- Step 2** In the **Organization VDCs** table, double-click the Org VDC where you want to create the internal network. The screen is refreshed with information about the selected VDC.
- Step 3** In the **Org VDC Networks** tab, in the toolbar, choose **Actions > Add Network**.



The **New Organization Network** wizard opens, guiding you through the configuration process.

Step 4 In the **Select Network Type** screen:

- a) Choose **Create a routed network by connecting to an existing edge gateway**.
- b) Choose the vShield Edge Gateway that you created in [Adding a vShield Edge Gateway on an Org VDC](#), on page 52.



Step 5 In the **Configure Network** screen:

- a) Enter the following information:
 - Gateway IP address
 - Network mask
 - DNS server IP address

- b) In the Static IP pool area, enter an IP address or an IP address range and click **Add**.
- Step 6** In the **Name and Description** screen, enter a name and description (optional) for the Org VDC internal network.
- Step 7** In the **Ready to Complete** screen, review the information for accuracy and click **Finish**.
-

Creating a Catalog

A catalog enables you to upload images from Cisco ICFPP to VCD.

For additional information about creating catalogs and selecting options, see your VMware vCloud Director documentation.

Before You Begin

Procedure

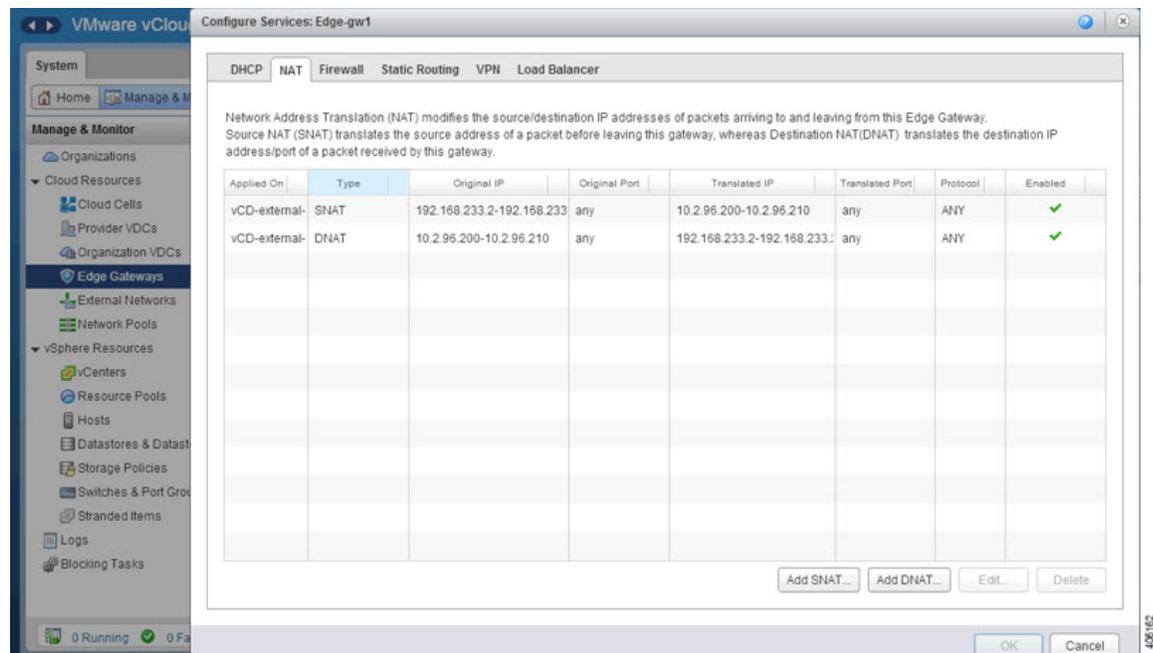
- Step 1** In the VCD GUI, choose **System > Manage & Monitor > Cloud Resources > Organization VDCs**.
- Step 2** In the **Organization VDCs** table, double-click the Org VDC in which to add the catalog. The screen is refreshed with information about the selected VDC.
- Step 3** Choose the **Catalogs** tab and, in the toolbar, choose **Actions > Add Catalog**. A dialog box opens with multiple tabs so that you can configure the catalog and user access.
- Step 4** In the **General** tab, enter a name and a description (optional) for the catalog.
- Step 5** In the **Sharing** tab:
- Click **Add Members**.
 - Choose the users or groups of users who can access the catalog.
 - In the **Access Level** field, choose the level of access for each user or group of users: Read-only, Read/Write, or Full Control.
- Step 6** In the **Storage** tab, choose the type of storage.
- Step 7** In the **Metadata** tab:
- From the **Type** drop-down list, choose the metadata type.
 - In the **Name** field, enter a name for this metadata entry.
 - In the **User access of metadata** field, choose the level of access for the metadata: Read/Write, Read-only, or Hidden.
 - In the **Value** field, enter a text value for the metadata entry.
- Step 8** After you have configured the catalog, click **OK**.
-

Verifying NAT and Firewall Service Configuration

When VCD is integrated with Cisco ICFPP, NAT and firewall services are configured automatically, enabling the vShield Edge Gateway to communicate with the external network. This procedure enables you to confirm that NAT and firewall services have been configured on the vShield Edge Gateway as expected.

Procedure

- Step 1** In the VCD GUI, choose **System > Manage & Monitor > Cloud Resources > Organization VDCs**.
- Step 2** In the **Organization VDCs** table, double-click the Org VDC where you created the vShield Edge Gateway ([Adding a vShield Edge Gateway on an Org VDC, on page 52](#)).
The screen is refreshed with information about the selected VDC.
- Step 3** In the **Edge Gateways** tab, right-click the required edge gateway and choose **Edge Gateway Services**.
- Step 4** In the Configure Services dialog box, confirm the following:
 - a) In the NAT tab, confirm that Source NAT and Destination NAT rules are displayed, as shown in the following example:



- b) In the Firewall tab, confirm that inbound traffic is allowed for the following destination ports and protocols:
 - 22—TCP
 - 443—TCP
 - 500—TCP, UDP
 - 4500—TCP, UDP
 - 6644—TCP, UDP

- 6646—TCP, UDP

The information should resemble the following example:

The screenshot shows the VMware vCloud Director interface for configuring services on 'ketan-EG'. The 'Firewall' tab is selected, and the 'Enable firewall' checkbox is checked. The default action is set to 'Deny'. Below this, a table lists 12 'InBound ACL Rules' with their respective configurations.

Rule Id	Name	Source	Destination	Protocol	Action	Log	Enabled
3	InBound ACL Rules	Any:Any	Any:22	TCP	Allow	-	✓
4	InBound ACL Rules	Any:Any	Any:6644	TCP	Allow	-	✓
5	InBound ACL Rules	Any:Any	Any:6644	UDP	Allow	-	✓
6	InBound ACL Rules	Any:Any	Any:6646	TCP	Allow	-	✓
7	InBound ACL Rules	Any:Any	Any:6646	UDP	Allow	-	✓
8	InBound ACL Rules	Any:Any	Any:443	TCP	Allow	-	✓
9	InBound ACL Rules	Any:Any	Any:500	TCP	Allow	-	✓
10	InBound ACL Rules	Any:Any	Any:500	UDP	Allow	-	✓
11	InBound ACL Rules	Any:Any	Any:4500	TCP	Allow	-	✓
12	InBound ACL Rules	Any:Any	Any:4500	UDP	Allow	-	✓



Additional Information

- [Related Documentation for Cisco Intercloud Fabric Provider Platform](#), page 59
- [Obtaining Documentation and Submitting a Service Request](#), page 60
- [Documentation Feedback](#), page 60

Related Documentation for Cisco Intercloud Fabric Provider Platform

The documentation listed below is available for Cisco Intercloud Fabric Provider Platform at the following URL:

<http://www.cisco.com/c/en/us/support/cloud-systems-management/intercloud-fabric/tsd-products-support-series-home.html>

General Information

Cisco Intercloud Fabric Provider Platform Release Notes

Install and Upgrade

Cisco Intercloud Fabric Provider Platform Installation Guide

Administration

Cisco Intercloud Fabric Provider Platform Administrator Guide

Troubleshooting and Alerts

Cisco Intercloud Fabric Provider Platform Troubleshooting Guide

Cisco Intercloud Fabric Documentation

The documentation listed below is available for Cisco Intercloud Fabric at the following URL:

<http://www.cisco.com/c/en/us/support/cloud-systems-management/intercloud-fabric/tsd-products-support-series-home.html>

Cisco Intercloud Fabric Release Notes
Cisco Intercloud Fabric Getting Started Guide
Cisco Intercloud Fabric Director REST API Guide
Cisco Intercloud Fabric Configuration Guide
Cisco Intercloud Fabric Firewall Configuration Guide
Cisco vPath and vServices Reference Guide for Intercloud Fabric
Cisco Intercloud Fabric User Guide
Cisco Intercloud Fabric Troubleshooting Guide

Cisco Nexus 1000V Documentation

[Cisco Nexus 1000V for VMware vSphere](#)
[Cisco Nexus 1000V for KVM](#)
[Cisco Nexus 1000V for Microsoft Hyper-V](#)

Cisco Virtual Security Gateway Documentation

[Cisco Virtual Security Gateway](#)

Cisco Cloud Services Router Documentation

[Cisco Cloud Services Router 1000V](#)

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation*, at: <http://www.cisco.com/c/en/us/td/docs/general/whatsnew/whatsnew.html>.

Subscribe to *What's New in Cisco Product Documentation*, which lists all new and revised Cisco technical documentation as an RSS feed and delivers content directly to your desktop using a reader application. The RSS feeds are a free service.

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to: intercloud-fabric-doc-feedback@cisco.com.

We appreciate your feedback.