

VIA User's Guide

VMware Cloud Foundation 2.3.1

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2015 - 2018 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About the VIA User's Guide 5

1 Overview 6

Cloud Foundation Deployment Options 7

Software Bundle 7

VIA Components 8

Database 8

Services 8

Components of a Physical Rack 9

2 Requirements for VIA 12

3 Setting up your Environment for Imaging 13

Rack Power 13

Network Cables 14

Rack Wiring 16

High level Wiring for Additional Racks 20

Detailed Wiring for Rack Components 21

Rack Component Ports 26

Physical Servers 29

BIOS Settings 30

Power Cycle Requirements 31

Network Switches 31

Laptop or Management Host 32

Virtual Machines 34

4 Pre-Imaging Checklist 36

5 Installing VIA 38

Installing VIA on a Laptop or Desktop 38

Installing VIA on a Management Host 39

6 Imaging Physical Racks 42

Image a Physical Rack 43

Upload Software Bundle 45

Add Server VIBs 46

Add Custom ESXi Installer ISO 47

Specify Imaging Details 48

- Monitor Imaging 50
- Verify Inventory 52
- Post Imaging Checks 53
- Download Inventory File 54
- Resume Imaging 55
- Image Additional Racks 57

- 7 Imaging Individual Devices 58**
 - Image Individual Server 58
 - Image New Management Switch 60

- 8 Review Alarms and Notifications 61**

- 9 Viewing the VIA Log File 62**

- 10 Viewing Results of an Imaging Run 63**
 - View Imaging History 63
 - View Inventory 64

- 11 Guidance on Server OOB Port Management 66**
 - OOB Ports Wiring to Management Switch 66
 - OOB Ports Configuration 66

- 12 Troubleshooting VIA 67**
 - ESXi Server has Incorrect BIOS Settings 67
 - ESXi Server has Bad SD Card 67
 - Management Switch Boots into EFI Shell 68
 - A switch did not start imaging 68
 - One or More Servers Are Not Discovered 69
 - Server Imaging Failed at Kickstart Delivery Task 70
 - Server Imaging Failed at Check Host Status Task 70

- 13 Cloud Foundation Glossary 71**

About the VIA User's Guide

The *VIA User's Guide* provides information about how to install VIA, manage software bundles, and image physical racks.

Intended Audience

This information is intended for anyone who wants to install or upgrade VIA and image physical racks. The information is written for experienced Windows or Linux system administrators who are familiar with virtual machine technology and datacenter operations.

Related Publications

The *VMware Cloud Foundation Overview and Bring-Up Guide* contains detailed information about the Cloud Foundation product, its components, and the network topology of an Cloud Foundation installation.

The *Administering VMware Cloud Foundation* provides information about how to manage a VMware Cloud Foundation™ system, including managing the system's physical and logical resources, managing users, configuring and deploying service offerings, and upgrading and monitoring the system.

VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <http://www.vmware.com/support/pubs>.

Overview

1

VMware Cloud Foundation is the unified SDDC platform that brings together vSphere, vSAN, NSX, and vRealize Suite into a natively integrated stack to deliver enterprise-ready cloud infrastructure for the private cloud.

The first step in deploying the Cloud Foundation solution is to image your hardware using VIA. VIA is a virtual appliance that installs VMware software on the physical racks in your datacenter using a collection of HTTP services.

A physical rack consists of a management switch, two Top of Rack (ToR) switches, and 4 to 32 physical servers. If you have multiple physical racks in your datacenter, the second rack must contain two inter-rack switches for inter-rack connectivity.

For imaging the rack, VMware provides the VIA OVF template and the Cloud Foundation software bundle. The software bundle contains the SDDC components.

The imaging infrastructure at the customer or partner site includes a laptop, desktop, or ESXi host (referred to as the management host) and a managed switch. The VIA OVA template is deployed on the laptop or management host and the software bundle is uploaded on the VIA VM. The laptop or management host is connected to the corporate network and to the private network used by the VIA VM. You use a browser (on the laptop) or jump VM (on the management host) to connect to the VIA VM and image the physical rack.

Figure 1-1. VIA Deployment

During imaging,

- The management, ToR, and inter-rack switches (if applicable) are configured.
- VMware ESXi is installed on each server in the physical rack.
- The Cloud Foundation software bundle, SDDC Manager VM, and SDDC Manager Utility VM, are installed on the primary host.

VIA runs an internal DHCP service to allocate private IP addresses to the switches and servers during imaging. After imaging is complete, VIA compiles a manifest file that provides an inventory of the physical rack components. The rack is now ready to be configured for Cloud Foundation, a process that is called bring-up. For more information on bring-up, see *VMware Cloud Foundation Overview and Bring-up Guide*.

This chapter includes the following topics:

- [Cloud Foundation Deployment Options](#)
- [Software Bundle](#)
- [VIA Components](#)
- [Components of a Physical Rack](#)

Cloud Foundation Deployment Options

Cloud Foundation provides flexibility in choosing on-premises deployment options.

Customers begin by sizing their Cloud Foundation deployment to determine the number of physical servers in their rack and number of racks. Each rack requires a minimum of 4 servers.

Customers have two deployment options for Cloud Foundation

- Deploy the Cloud Foundation software on qualified vSAN ReadyNodes in your datacenter.
Customers can start with qualified hardware (qualified ReadyNodes and qualified switches) in their datacenter, wire it up, and deploy the Cloud Foundation software stack on the ready system. For information on qualified hardware, see [VMware Compatibility Guide](#).
- Purchase a fully integrated system that combines software and hardware from select VMware partners.
The partner works with a VMware representative to complete the Site Readiness document. This translates into a bill of materials (BoM) consisting of both hardware and software components. With this BoM in hand, the partner assembles the rack and images it. The partner then ships the system, consisting of physical racks, servers, server sub-components, power distribution units, switching infrastructure and the Cloud Foundation software, to customers.



Deploying VMware Cloud Foundation on Qualified Hardware
(http://link.brightcove.com/services/player/bcpid2296383276001?bctid=ref:video_deploy_cloud_foundation_hardware)

Software Bundle

The software bundle is a collection of all the software, configuration files, utilities, and tools used by VIA to image a physical rack. It contains a manifest file that lists the contents of the bundle. The bundle is based on a hardware bill-of-materials (BoM), that includes specific servers, switch models, and their component level configurations.

In addition to the core bundle, you can upload 3rd party VIBs, and custom ESXi installer ISOs on VIA.

The core software bundle contains the following software:

- vSphere (vCenter Server and ESXi)

- NSX
- vSAN
- vRealize Log Insight
- SDDC Manager
- Platform Services Controller
- VMware Horizon
- App Volumes
- vRealize Operations
- vRealize Automation

See the *VMware Cloud Foundation Release Notes* for the software component versions.

VIA Components

VIA uses multiple components to track and perform the imaging process. This section describes these components, but you do not need to perform any configuration on them.

Database

VIA stores information about all activities during an imaging run in an MongoDB database. This includes current imaging information as well as the previous imaging status. All entities utilized by the imaging process are stored as an entry in the database. These entities include the software bundle, imaged component, manifests, user information, and hardware information.

Services

In order to handle disparate requests that may be required to service its components, VIA deploys multiple services. Each service has a specific goal, and is instantiated based on the state of the imaging activity.

Switch Service

A Switch Service is developed for each switch vendor by using an imaging developer kit. The developer kit provides the imaging orchestration engine, API, data models, and extension points to create imaging tasks (PRE, INTRA, POST, INVENTORY), custom controller, and automated integration tests. VIA loads these services on demand and discovers the switch type (management, ToR, inter-rack) supported by these services.

ESXi Service

The ESXi Service images the servers in the physical rack. It uses the same developer kit as the Switch Service.

Core Platform Service

Core Platform is the main VIA web service which supports external facing API, DHCP, bundle management, imaging workflow orchestration, and UI. This service automatically loads the Switch and ESXi services on demand based on the BOM and bundle and then orchestrates the imaging workflow. This service can image an entire rack as well as individual devices.

Components of a Physical Rack

VMware recommends that you use a white cabinet that is 19" wide with 42 Rack Units (RU) for the physical rack. The cabinet must have a loading capacity of 2000 lbs and have adjustable levelling feet with heavy duty casters and seismic bracing. Since switches do not cover the full shelves, the cabinet must have a grill on one side for proper airflow.

Table 1-1. Rack Components

Component	Rack 1	Additional Racks
PDUs	4	4
Console serial switch	1	1
Inter-rack switches	NA	2 (Rack 2 only)
TOR Switches	2	2
Management switch	1	1
Servers	Up to 32	Up to 32

■ PDUs

Each physical rack must have 4 PDUs (2 primary and 2 standby) even if it contains less than 32 servers. It is recommended that the primary PDUs be blue and the standby be red. The primary PDUs must be placed on the rear left side and the standby PDUs must be placed on the rear right side of the cabinet. The capacity requirements for each PDU are:

- 208 V
- 30 AMP
- 3 phase
- 60 Hz/50 H

The plug type needs to be determined based on the customer's environment.

■ Console serial switch

Each physical rack contains a 16-port console serial switch. The console serial switch is connected to all the other switches in the rack and is used for troubleshooting.

■ Inter-rack switches

Rack 2 in your Cloud Foundation system contains two 32 x 40 GE inter-rack switches. These switches connect multiple racks by using uplinks from the Top of Rack switches.

Inter-rack switches must not be connected to a corporate network. They are only used for ToR connectivity between physical racks.

- Top of Rack (ToR) switches

Each rack contains two 1RU 48-port 10 GE ToR switches with four 40 GE uplinks. Servers in each rack are connected to both ToRs. The ToRs on rack 1 connect Cloud Foundation to the corporate network.

- Management switch

Each rack contains a 1 GE management switch, which is used to access the physical switches. The management switch is connected to the management ports of the ToR and inter-rack switches (on Rack 2 only). The management switch is also connected to the data ports of the ToR switches. VIA uses the management port connection to the ToR switch for imaging and configuration of the management switch, and the data port connection for imaging and configuration of the other switches and servers.

- Servers

A Cloud Foundation rack can contain 4 to 32 heterogeneous servers. Servers must be from the same vendor, but can be of different models and sizes with variable CPU, memory, storage size and type, and disk configurations (e.g. hybrid or all flash).

You can select the servers you want to use for the management domain as well as workload domains. This gives you the flexibility to use say hybrid servers for the management domain and all flash servers for workload domains.

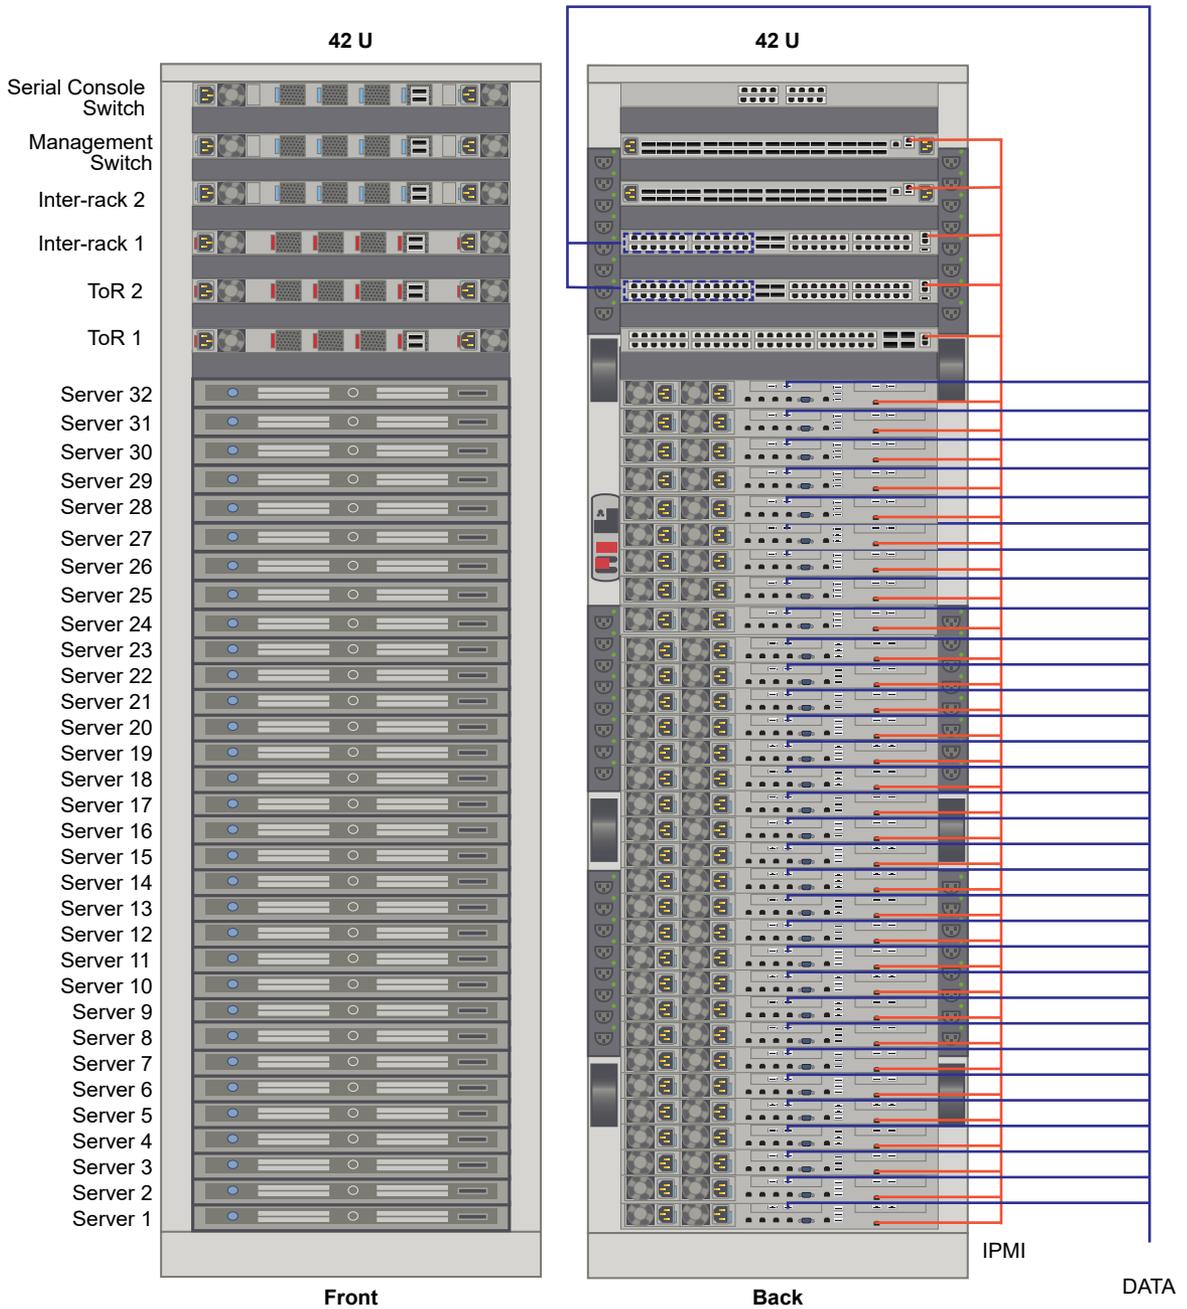
Table 1-2. Server Configuration for Cloud Foundation

Component	Minimum	Maximum
CPU per server	Dual-socket, 8 cores per socket	Dual-socket, no maximum on cores per socket
Memory per server	256 GB	1.5 TB
Storage per server	4 TB for capacity tier. Follow vSAN guidelines for cache tier sizing as described in VMware vSAN Design and Sizing Guide . Each server to be used in the domain must contain at least 3 capacity tier disks.	8 disks per disk-group and 2 disk-groups per host. Follow vSAN guidelines for cache tier sizing as described in VMware vSAN Design and Sizing Guide .
	Note Cloud Foundation only supports vSAN RAID controllers in pass-through mode.	
NICs per server	Two 10 GbE NICs and one 1 GbE BMC NIC	Two 10 GbE NICs and one 1 GbE BMC NIC

Table 1-2. Server Configuration for Cloud Foundation (continued)

Component	Minimum	Maximum
Servers per rack	Four 1U or 2U servers A minimum of 7 servers are required for VI and VDI workload creation.	32 1U servers or 16 2U servers
Rack	1	8

Figure 1-2. Example Physical Rack Configuration



Requirements for VIA

2

VIA requires the following infrastructure.

- You need a laptop/desktop or an ESXi host (called management host) to run the VIA VM. The laptop need to be connected to the management switch of the rack being imaged. If there are multiple racks to be imaged, this would mean physically moving the laptop for imaging each rack. The management host is connected to all the management switches through an internal switch, so the connection to the rack being imaged can be managed remotely.
- Desktop or laptop (Windows or Mac) with 4 GB memory and a multi core processor to access the jump VM. A Windows laptop must have Workstation 8 or later and a Mac should have Fusion 4 or later installed on it. You also need a network adapter, a cable, and a minimum 4-port unmanaged switch.
- Management host - a standalone VMware vSphere ESXi 6.0 or later server to host the Windows jump VM. The management host must have at least two NICs, with one NIC connected to the corporate network and one NIC connected to the private network.
- If you are using a management host for imaging, you need a jump VM to access VIA
- Supported 1GE private managed switch with RJ45 ports and Cat 5/5E cables and at least 3 ports. Private indicates that it is being used only by you. The managed switch provides the ability to configure, manage, and monitor your LAN, which gives you greater control over how data travels over the network and who has access to it.

Setting up your Environment for Imaging

3

You must inspect the components of the physical rack, verify cable connectivity, and validate BIOS settings before beginning the imaging process.

If you are imaging your own rack, refer to the [VMware Compatibility Guide](#) to ensure that the hardware in your datacenter is supported for Cloud Foundation.

If you are a partner imaging a rack for your customer, review the VMware Bill Of Materials (BOM) before you set up your environment for imaging.

This chapter includes the following topics:

- [Rack Power](#)
- [Network Cables](#)
- [Rack Wiring](#)
- [Rack Component Ports](#)
- [Physical Servers](#)
- [Network Switches](#)
- [Laptop or Management Host](#)
- [Virtual Machines](#)

Rack Power

Ensure that rack power meets the following requirements.

- Verify that each device in the rack has a connection to each PDU.
- VMware recommends that you cable each server to the nearest power port so that the cable length can be kept to a minimum. Length of power cables should be as follows.
 - From the Physical Server: (9) .5m (3) 1m
 - From the Top-of-Rack Switch: 1.5m
 - From the Inter-rack Switch: .5m

It is common for power cables within a rack to be longer than required. However, if excess cabling is not managed properly, it may create electromagnetic interference. Avoid bundling of excess cables as this may lead to the cables being damaged due to bending.

- The power connector from the PDU must match the power connector in the Site Readiness Assessment.
- Power cables must be seated properly from each device to the PDU.
- The cables connect the primary PDUs to the other components must be blue and the cables from the secondary PDUs must be red.
- Power cables should not be in an area where there is a risk of touching sharp edges, excessive heat, or subject to pinching between sliding rails.

Network Cables

Proper management of network cables promotes the elimination of crosstalk and interference, cooler performance, improved maintenance, and easier upgrades. Incorrect cable management may result in damage or failure, which may lead to data transmission errors, performance issues, or system downtime. This section contains cable color and management recommendations. You can adapt the recommendations to suit your environment.

Regardless of the number of servers in each rack, cables must be in place for 32 servers. Ideally, data and power cables must be at opposite ends of the physical rack. If they are aggregated in a bundle or run parallel to each other, induction may introduce electromagnetic interference.

Cable Colors

Using specific colors for cables from each device makes for easier troubleshooting.

- All cables from the management switch (except those going to the ToRs): yellow
- Management switch ports 49 and 50 going to the ToRs: black
- ToR 1 cables to servers: blue
- ToR 2 cables to servers: red
- ToR 1 and ToR 2 connections to inter-rack switches: orange
- Console serial switch connections: grey

Cable Type and Length

The Telecommunications Industry Association (TIA) and the Electronic Industries association (EIA) structured cabling standards define how to design, build, and manage cabling systems. The specification is TIA/EIA-568-A. When used for 10/100/1000BASE-T Category 6 (Cat 6) cable length can be up to 100 meters (328 ft). This distance includes up to 90 meters (295 ft) of

horizontal cabling between the patch panel and the wall jack, and up to 10 meters (33 ft) of patch cabling. When used for 10GBASE-T, Cat 6 cable length is reduced to 55 meters (180 ft) assuming minimal exposure to crosstalk. Category 6A (Cat 6A) does not have this limitation and can run at the same distances as 10/100/1000BASE-T.

Cable Bend Radius

Modifying the geometry of a cable can impair data transmission and affect performance. When a cable is tied or tightly looped, the pairs within the cable jacket can be separated impacting the integrity of the cable. Therefore, bend radius should be considered when verifying cable management.

- The minimum bend radius of a twisted pair patch cable is 4x the external cable diameter, and the minimum bend radius of an LC-type fiber optic cable is 0.8" (~2cm) and SC-type fiber optic cable is 1" (~3cm).
- Where articulated arms or rail slides are used, there must be sufficient slack in the cable to allow operation.
- No creases in the sheathing should be visible on any cable.

Cable Routing

Improperly routed cables can contribute to thermal issues, make field replaceable units difficult to access, or impact performance.

Cable ties can damage cables due to excessive over tightening or by violating the bend radius of a cable. Cable ties also increase service time when an add, move, or change request is received. Cables should be bundled with Velcro straps where possible to avoid damage, simplify addition or removal of cables, and reduce service times.

- Use velcro straps instead of cable ties.
- Network cables should not be in an area where there is a chance of contacting sharp edges, excessive heat, or subject to pinching between sliding rails.
- Cables must be free of tension. Where articulated arms or rail slides are used, there must be sufficient slack in the cable to prevent the cables from being stressed.
- Forced air cooling is recommended to draw cool air from the front of the rack and push warm air out the back.
- Ventilation slots, power supplies, and rear fans must be clear of cable obstructions.
- Field replaceable units such as power supplies must be clear of any cable obstructions that may prevent access for service.

Cable Labeling

Partners must label the cables in their datacenter. Properly labeled cables reduce troubleshooting time since it is easier to trace and validate connections.

Cable Testing

Cable testing ensures that the installed cabling links provide the transmission capability to support the data communication required.

Several tools are available for copper testing. Tests fall into three categories: Verification, Qualification, and Certification. Verification tools are used to perform basic continuity, cable length, and open connection verification. Qualification tools can provide information that details the cable capabilities, e.g. supports 10GBase-T. Certification tools determine whether the cable meets TIA standards such as TIA-568-B.

Options for testing SFP+ and QSFP+ cables are limited. Because handheld cable testers are not available, many network administrators typically reserve ports between two adjacent switches, then connect a suspect cable between active ports to determine if the cable is functional.

- Review the test print out to confirm that the cables passed the test.
- Cables from the physical server 10G interface to the ToR switches must be tested prior to installation. They must be seated properly.
- Each 10G interface must be connected to a separate ToR switch.
- Inter-switch SFP+ and QSFP+ cables must be tested prior to installation.
- Each 40G QSFP+ cable from the ToR switch must be connected to a separate inter-rack switch.
- There must be two 40G QSFP+ cable connections between each ToR switch and inter-rack switch.
- Inter-switch SFP+ and QSFP+ cables must be seated properly.

Rack Wiring

Download VCF Wiremap from the Product Downloads page and connect the wires in your physical rack according to the wiremap. This section contains the logical views of the wiremaps.

High Level Wiring for Rack with Dell Management Switch

Figure 3-1. Wiremap for rack 1 with Cisco ToR Switches and Dell Management Switch

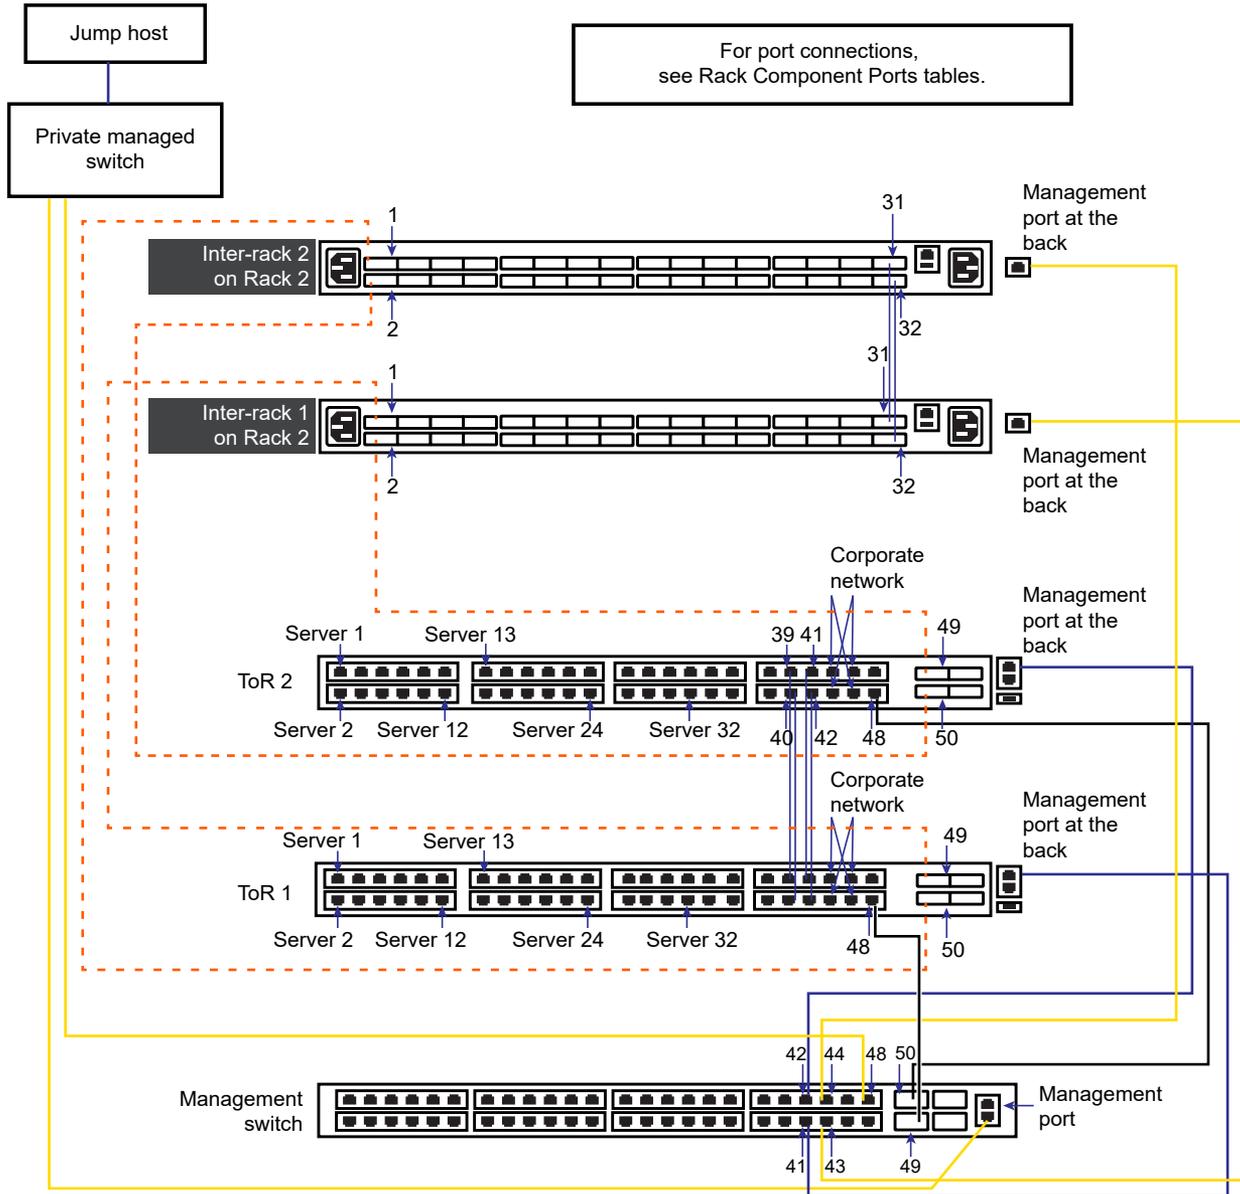
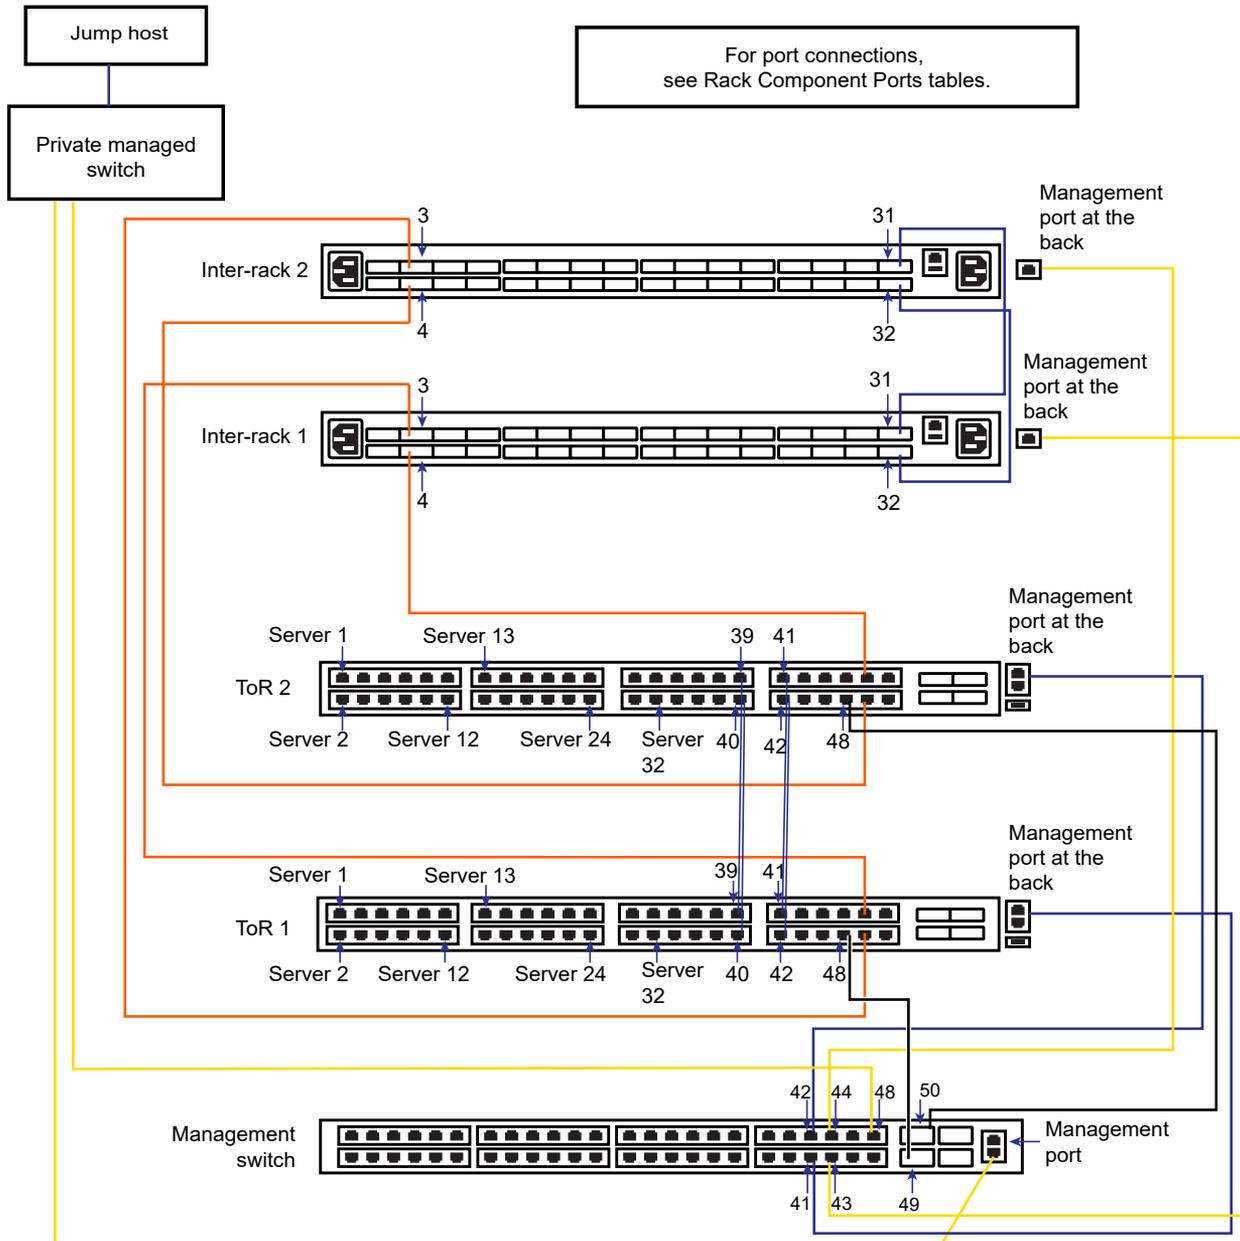


Figure 3-2. Wiremap for rack 2 with Cisco ToR Switches and Dell Management Switch



High Level Wiring for Rack with Quanta Management Switch

Figure 3-3. Wiremap for rack 1 with Cisco ToR Switches and Quanta Management Switch

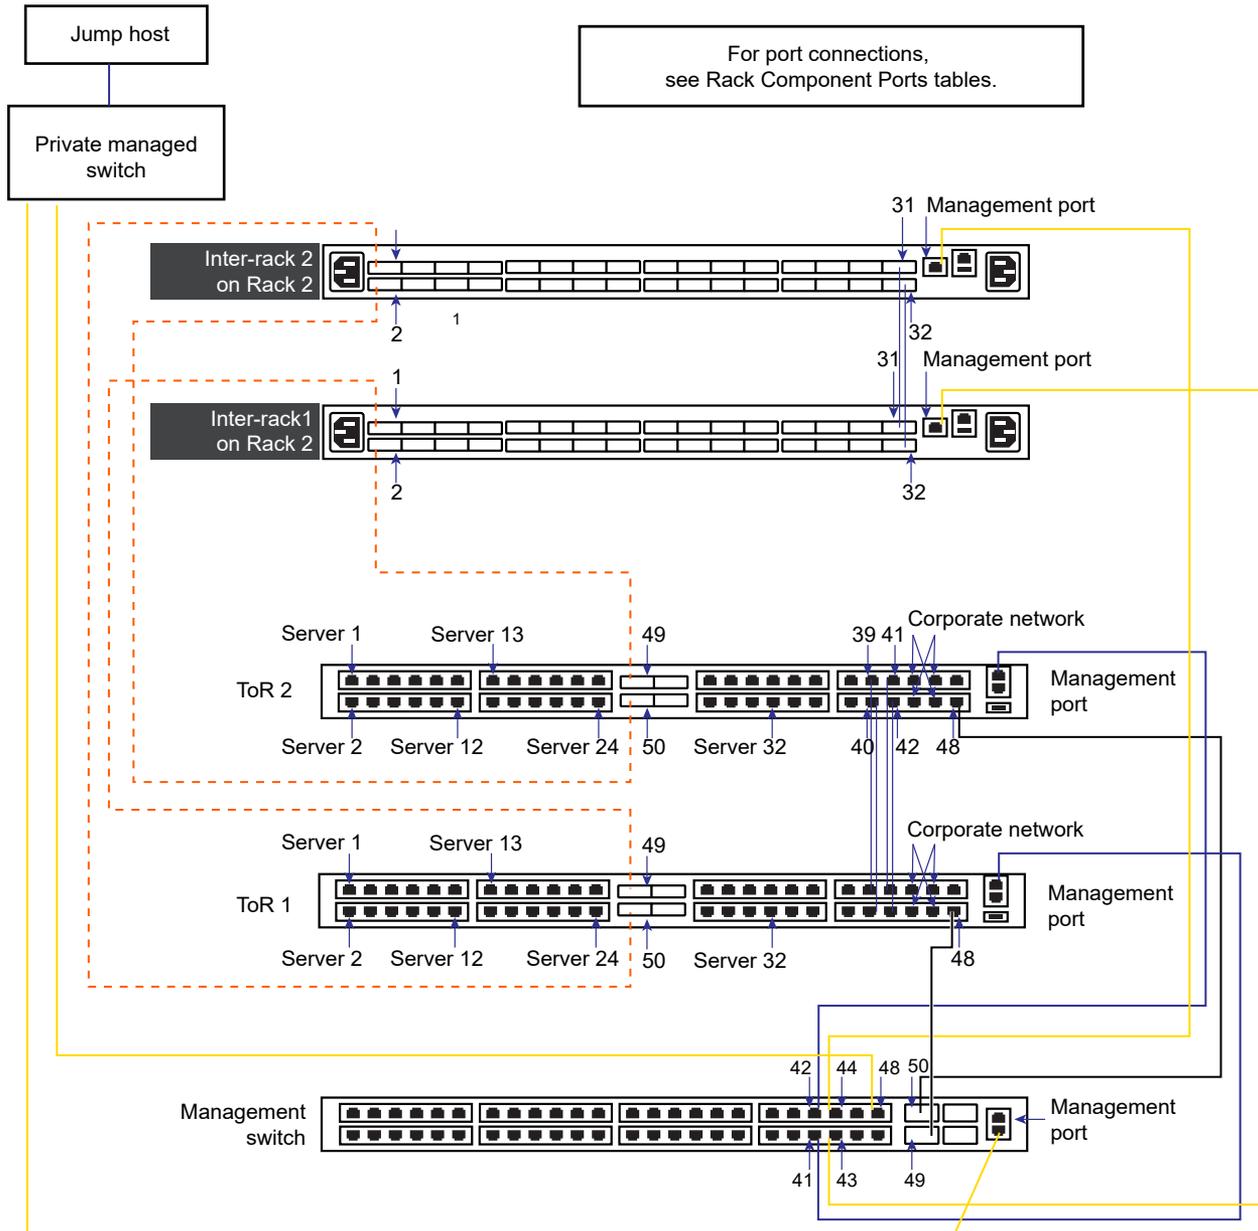
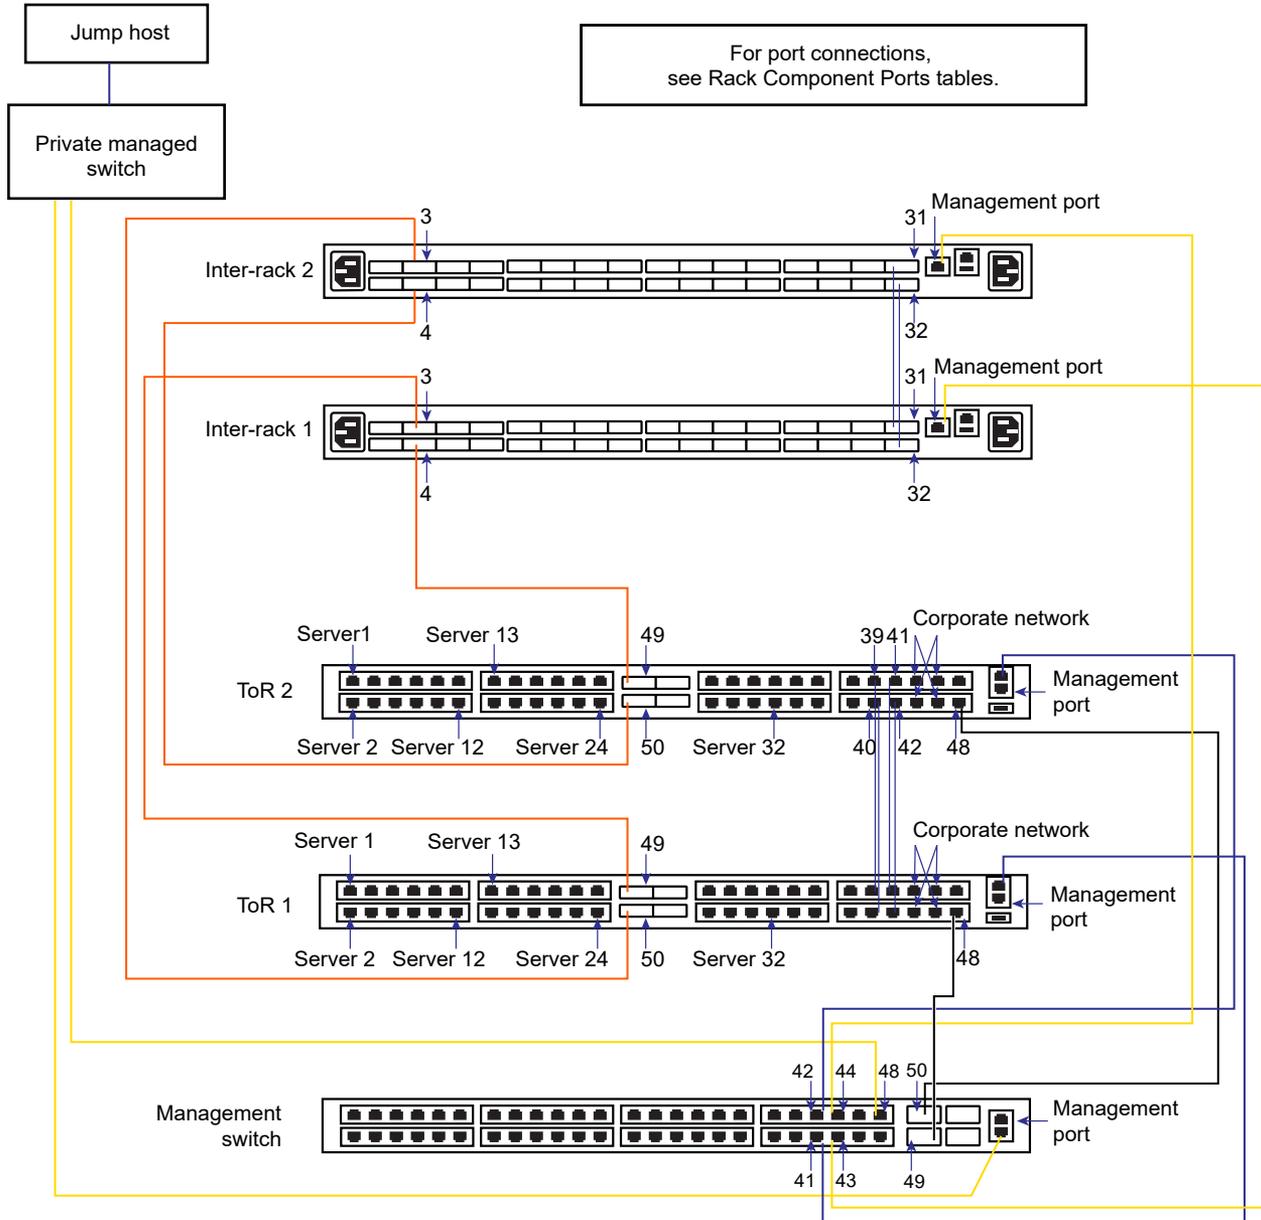


Figure 3-4. Wiremap for rack 2 with Cisco ToR Switches and Quanta Management Switch



High level Wiring for Additional Racks

Rack 2 in the integrated system powered by Cloud Foundation must include two inter-rack switches for inter-rack connectivity. The inter-rack switches are connected during the physical environment inspection, but must be disconnected before imaging the rack.

Additional physical racks do not contain inter-rack switches. ToR switches in the additional physical racks are connected to the two inter-rack switches in rack 2.

Detailed Wiring for Rack Components

This section shows the wiring for each component in a physical rack. These are example sketches. Port placement and numbering on components in your environment may be slightly different.

Figure 3-5. Management Switch

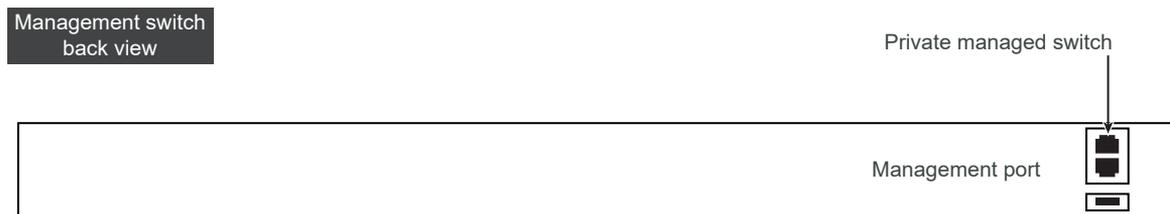
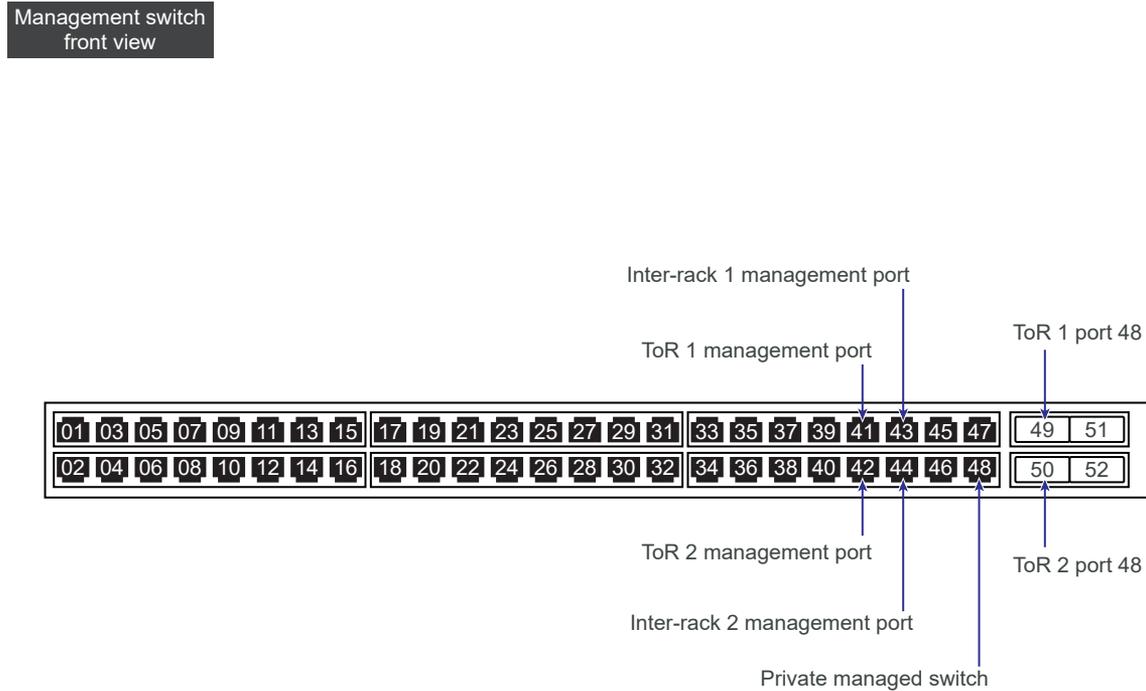


Figure 3-6. Inter-rack Switch 2

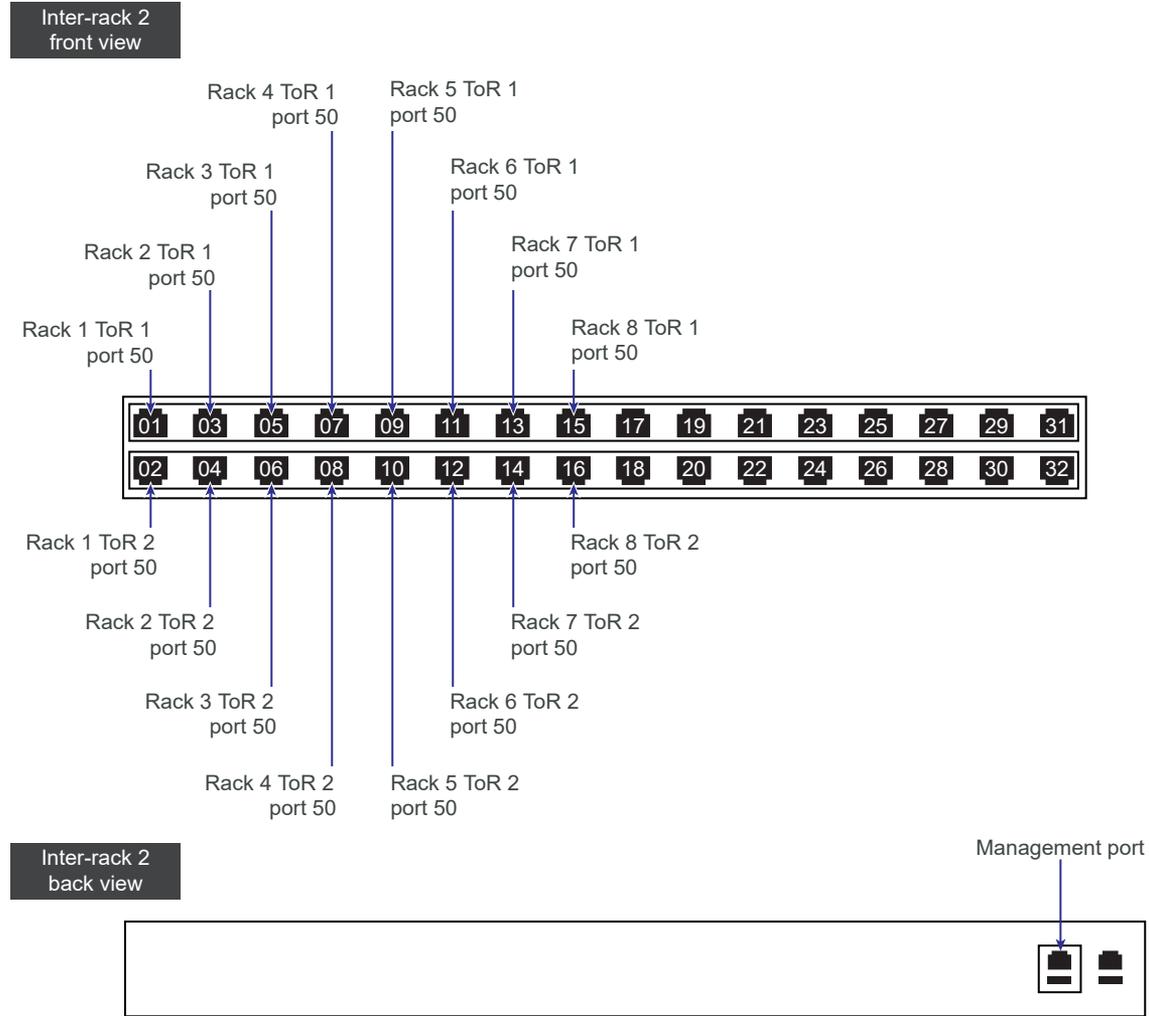


Figure 3-7. Inter-rack Switch 1

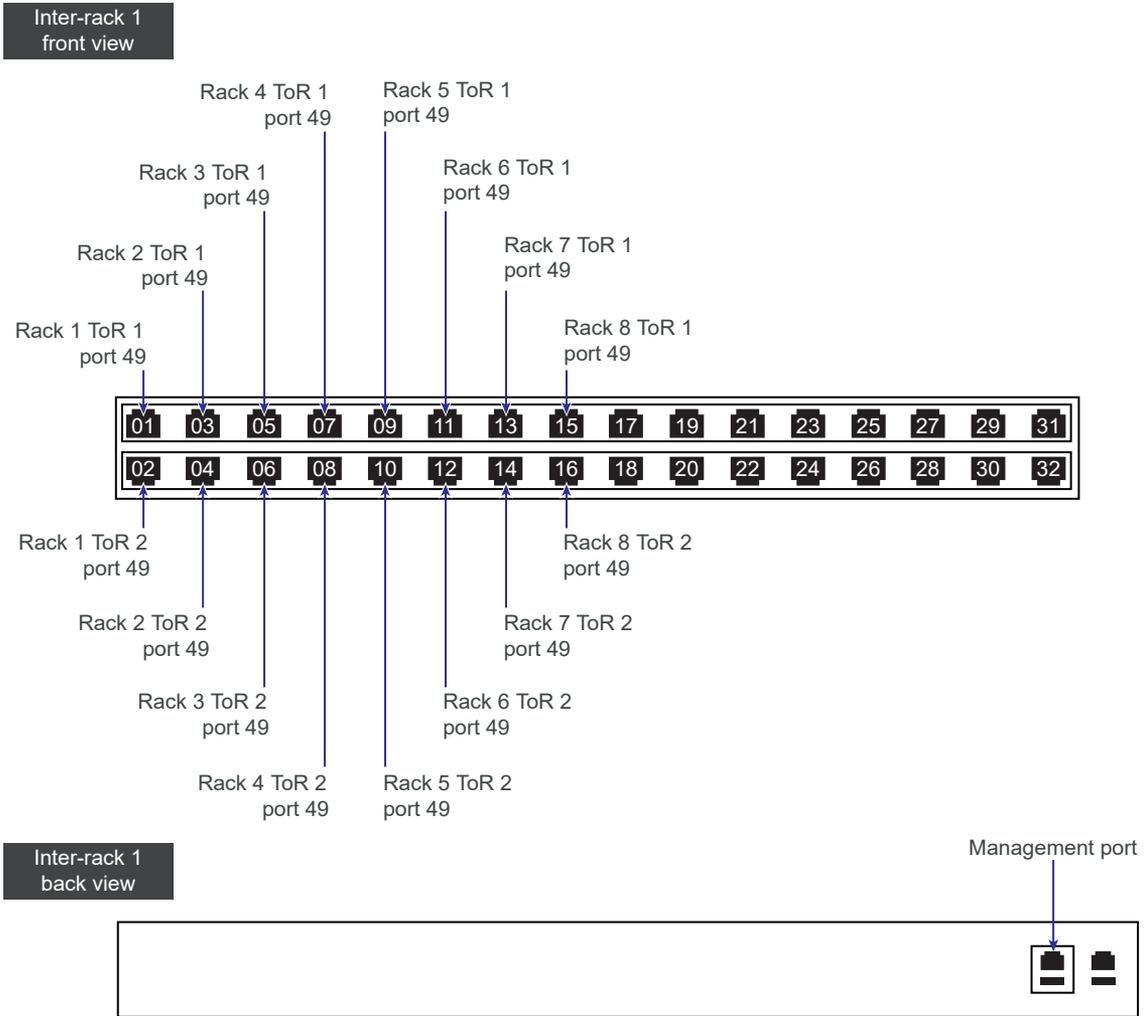


Figure 3-8. ToR 2

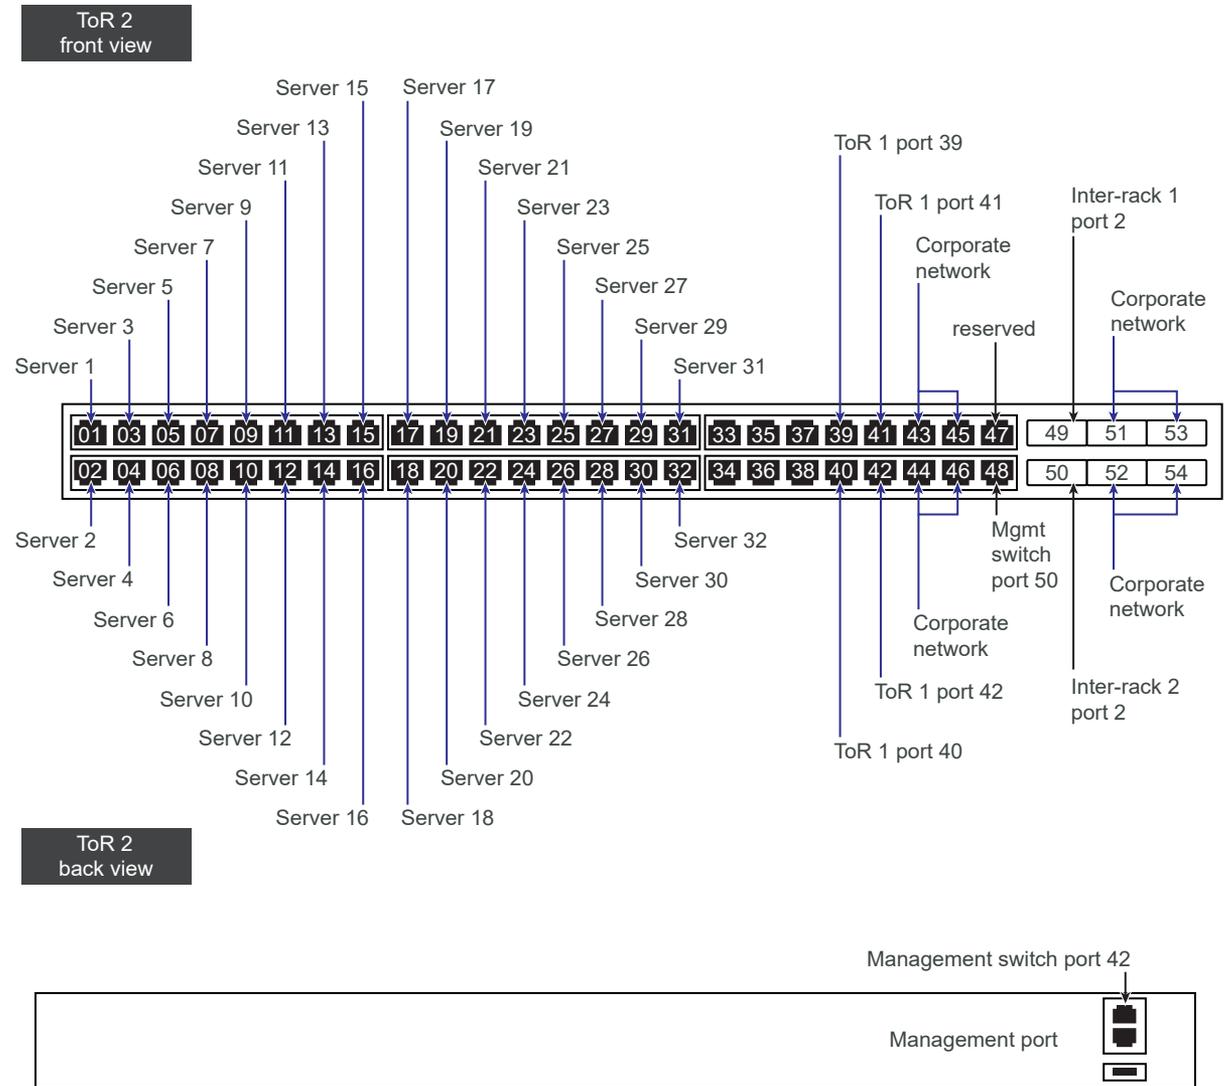


Figure 3-9. ToR 1

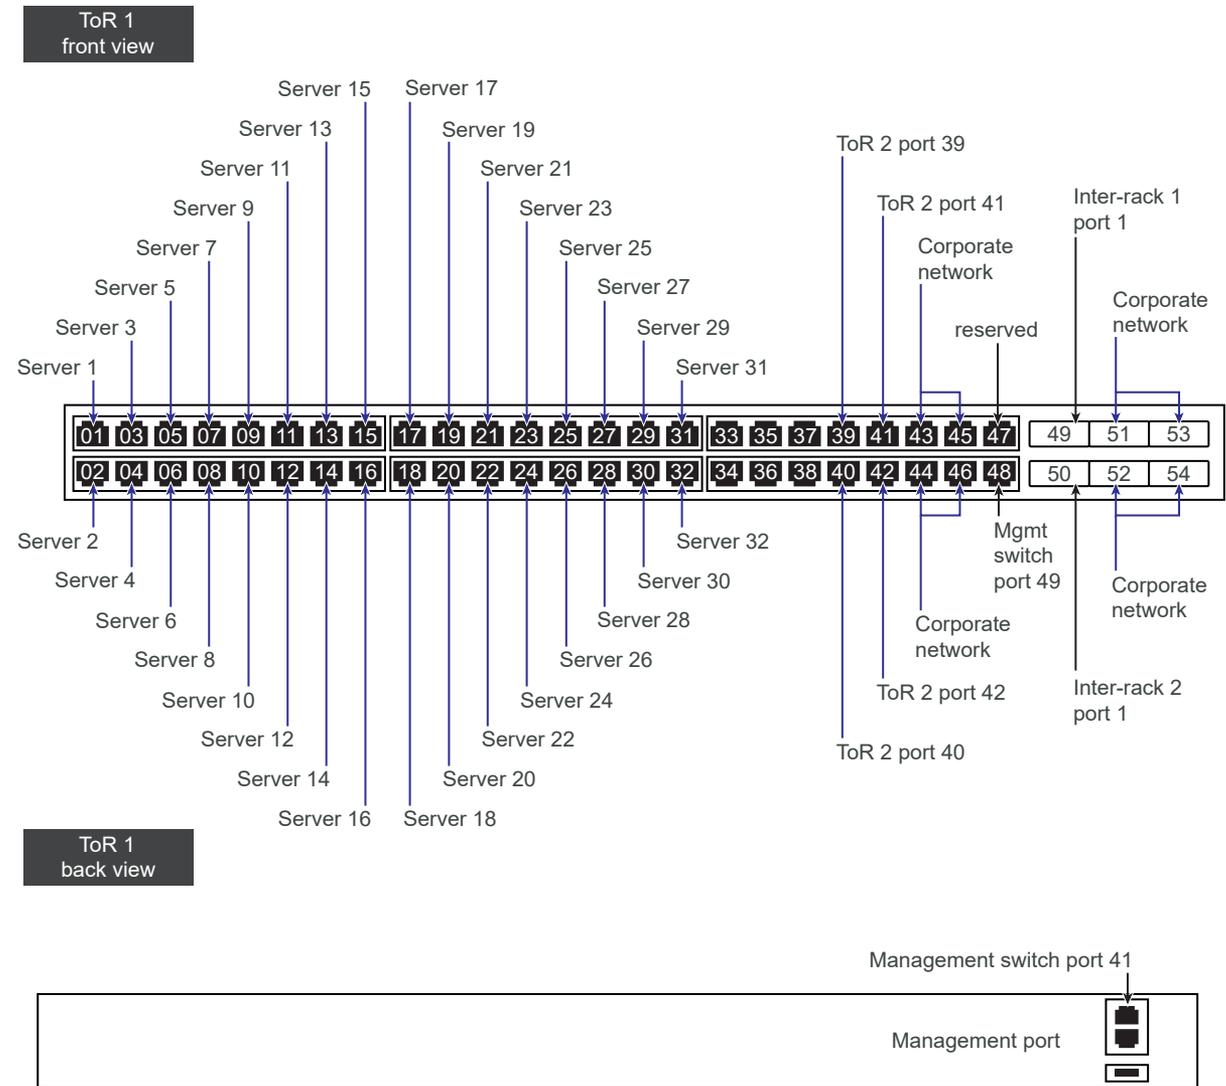
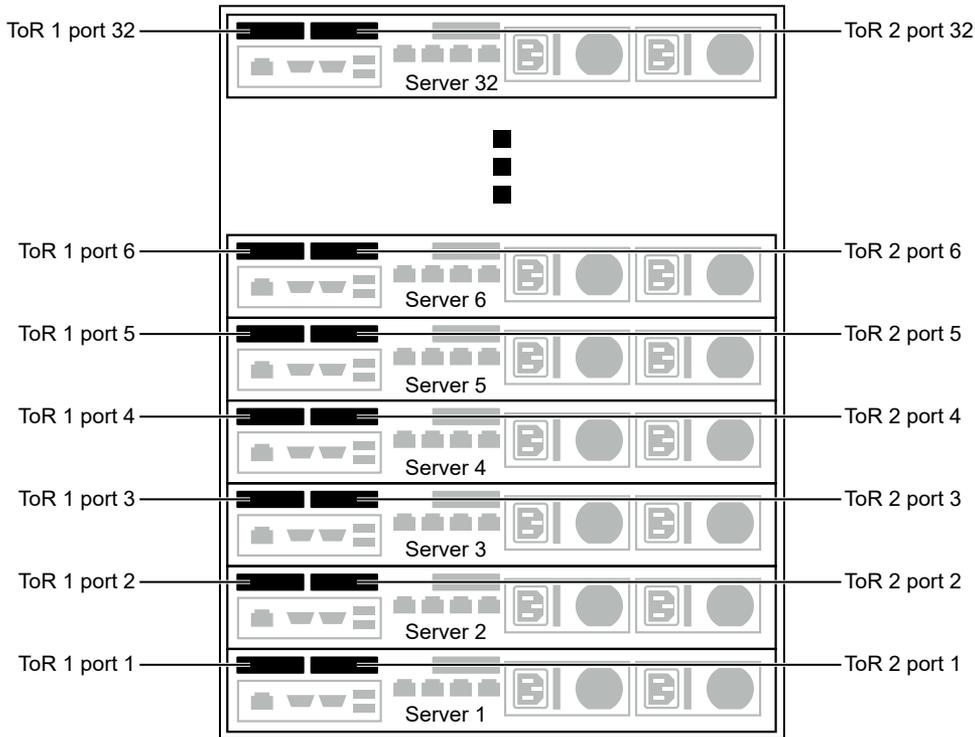


Figure 3-10. Servers



Rack Component Ports

Refer to the tables below for port connectivity information using Cisco 9372PX as the illustrative example. Connections in your environment may vary based on the actual switches being used.

Console Serial Switch

Port Number	Connects To
1	Management switch console port
2	ToR 1 console port7
3	ToR 2 console port
4	Inter-rack 1 console port
5	Inter-rack 2 console port
6	PDU 1
7	PDU 2
8	PDU 3
9	PDU 4
10 - 16	Not connected

Inter-rack 2 (Rack 2 only)

Port Number	Speed	Connects To
1	40 Gbps	Rack 1 ToR 1 port 50
2	40 Gbps	Rack 1 ToR 2 port 50
3	40 Gbps	Rack 2 ToR 1 port 50
4	40 Gbps	Rack 2 ToR 2 port 50
5	40 Gbps	Rack 3 ToR 1 port 50
6	40 Gbps	Rack 3 ToR 2 port 50
7	40 Gbps	Rack 4 ToR 1 port 50
8	40 Gbps	Rack 4 ToR 2 port 50
9	40 Gbps	Rack 5 ToR 1 port 50
10	40 Gbps	Rack 5 ToR 2 port 50
11	40 Gbps	Rack 6 ToR 1 port 50
12	40 Gbps	Rack 6 ToR 2 port 50
13	40 Gbps	Rack 7 ToR 1 port 50
14	40 Gbps	Rack 7 ToR 2 port 50
15	40 Gbps	Rack 8 ToR 1 port 50
16	40 Gbps	Rack 8 ToR 2 port 50

Inter-rack 1 (Rack 2 only)

Port Number	Speed	Connects To
1	40 Gbps	Rack 1 ToR 1 port 49
2	40 Gbps	Rack 1 ToR 2 port 49
3	40 Gbps	Rack 2 ToR 1 port 49
4	40 Gbps	Rack 2 ToR 2 port 49
5	40 Gbps	Rack 3 ToR 1 port 49
6	40 Gbps	Rack 3 ToR 2 port 49
7	40 Gbps	Rack 4 ToR 1 port 49
8	40 Gbps	Rack 4 ToR 2 port 49
9	40 Gbps	Rack 5 ToR 1 port 49
10	40 Gbps	Rack 5 ToR 2 port 49
11	40 Gbps	Rack 6 ToR 1 port 49
12	40 Gbps	Rack 6 ToR 2 port 49
13	40 Gbps	Rack 7 ToR 1 port 49
14	40 Gbps	Rack 7 ToR 2 port 49

Port Number	Speed	Connects To
15	40 Gbps	Rack 8 ToR 1 port 49
16	40 Gbps	Rack 8 ToR 2 port 49

ToR 2 (e.g. Cisco 9372PX)

Port Number	Speed	Connects To
1 - 32	10 Gbps	node 1 - node 32 where port 1 connects to node 1, port 2 connects to node 2, and so on
33 - 38	NA	Not connected
39 - 42	10 Gbps	ToR 1 ports 39 - 42
43 - 46	10 Gbps	Corporate network as required (see note below table)
47	Blank	
48	1Gbps	Management switch port 50
49	40 Gbps	Inter-rack 1 port 2
50	40 Gbps	Inter-rack 2 port 2
51 - 54	40 Gbps	Corporate network as required (see note below table)
Management	1 Gbps	Management switch port 42

Note Depending on the switches in your environment, connect two 40 Gbps ports or multiple 10 Gbps ports to your corporate network.

ToR 1 (e.g. Cisco 9372PX)

Port Number	Speed	Connects To
1 - 32	10 Gbps	Node 1 - node 32 where port 1 connects to node 1, port 2 connects to node 2, and so on
33 - 38	NA	Not connected
39 - 42	10 Gbps	ToR 2 ports 39 - 42
43 - 46	10 Gbps	Corporate network as required (see note below table)
47	Blank	
48	1Gbps	Management switch port 49
49	40 Gbps	Inter-rack 1 port 1
50	40 Gbps	Inter-rack 2 port 1
51 - 54	40 Gbps	Corporate network as required (see note below table)
Management	1 Gbps	Management switch port 41

Note Depending on the switches in your environment, connect two 40 Gbps ports or multiple 10 Gbps ports to your corporate network.

Management Switch

Port Number	Speed	Connects To
1 - 32	1 Gbps	Ports 1 - 32 are blank. You can connect these ports to the BMC ports on servers. This is optional and is not required for imaging. See "Guidance on Server OOB Port Management" in the <i>VIA User's Guide</i> .
33 - 40	NA	Not connected
41	1Gbps	ToR 1 management port
42	1Gbps	ToR 2 management port
43	1Gbps	Inter-rack 1 management port
44	1Gbps	Inter-rack 2 management port
45-47	NA	Not connected
48	1Gbps	Private managed switch
49	10 Gbps	ToR 1 port 48
50	10 Gbps	ToR 2 port 48
51-52	NA	Not connected
Management port		Private managed switch

Note PDU ports are not reflected in the table above.

Physical Servers

This section lists the recommended Rack Unit (RU) location of each device.

Hardware Devices

Table 3-1. Physical Device Location in Rack 1 and Rack 3 - Rack n

RU Location	Device
42	Console serial switch
41	Blank
40	Management switch
39	Blank
38	ToR 2
37	Blank
36	ToR 1
33-35	Blank
1-32	Nodes 1-32

Table 3-2. Physical Device Location in Rack 2

RU Location	Device
42	Console serial switch
41	Management switch
40	Inter-rack 2
39	Blank
38	Inter-rack 1
37	Blank
36	ToR 2
35	Blank
34	ToR1
33	Blank
1-32	Nodes 1-32

Power

All servers must have redundant power supplies and each power supply must be connected to a separate rack PDU.

Airflow

Install the servers to allow front-to-back airflow.

Firmware Settings

Ensure that the firmware settings are set correctly as per the BoM.

BIOS Settings

The BIOS settings for each server in the physical rack must match the values given below. It is a good practice to set the values in the order specified below to avoid any imaging failures.

If BIOS changes are needed to set the boot device, disable all NICs on each server except for the management port and dual port card. If the server can boot from the connected NIC cards, you do not need to disable the additional NICs.

For information on how to set the required BIOS values, refer to the vendor documentation.

Setting	Value
CPU and Performance Settings	Set per ESXi recommendations. See VMware vSphere documentation.
Onboard SATA controller	AHCI mode
Required only for servers with a SATADOM or M2 as boot disk	
Boot Order	Network First

Setting	Value
Second in Boot Order	Boot Disk
Boot mode	Legacy or UEFI
Storage Controller	Passthrough Mode

Reboot the servers after changing the BIOS settings to ensure that the changes are in place.

Power Cycle Requirements

VIA does not reboot servers during the imaging process. You must manually reboot the servers when the VIA UI displays a notification asking you to do so.

You can assign OOB addresses to the servers in your rack. For more information, see [Chapter 11 Guidance on Server OOB Port Management](#). Or you can manage the IP assignment in other ways.

Network Switches

Verify that SFP+ connectors are supported by the switches in your environment.

Power

- All switches must have redundant power supplies.
- Each power supply must be connected to a separate rack PDU.

Airflow

Switches must be installed to allow front-to-back airflow.

Switch Mode

Ensure that the mode for each switch is set such that it is ready to accept a new image and configuration.

Table 3-3. Required Mode for Supported Vendors

Vendor	Type of Switch	Mode
Arista	■ ToR	ZTP
	■ Inter-rack	
Dell EMC	■ ToR	BMP
	■ Inter-rack	
Cisco	■ ToR	POAP Recommended NXOS version is 7.0.3.I7.1.
	■ Inter-rack	
Dell EMC	Management	ONIE Recommended version is 3.24.1.2.

Laptop or Management Host

You need a laptop or management host where you install VIA.

Laptop

You need a desktop or laptop (Windows or Mac) with 4 GB memory, a multi core processor, and two NICs. A Windows laptop must have Workstation 8 or later and a Mac should have Fusion 4 or later installed on it. You also need a network adapter, a cable, and a 4-port unmanaged switch.

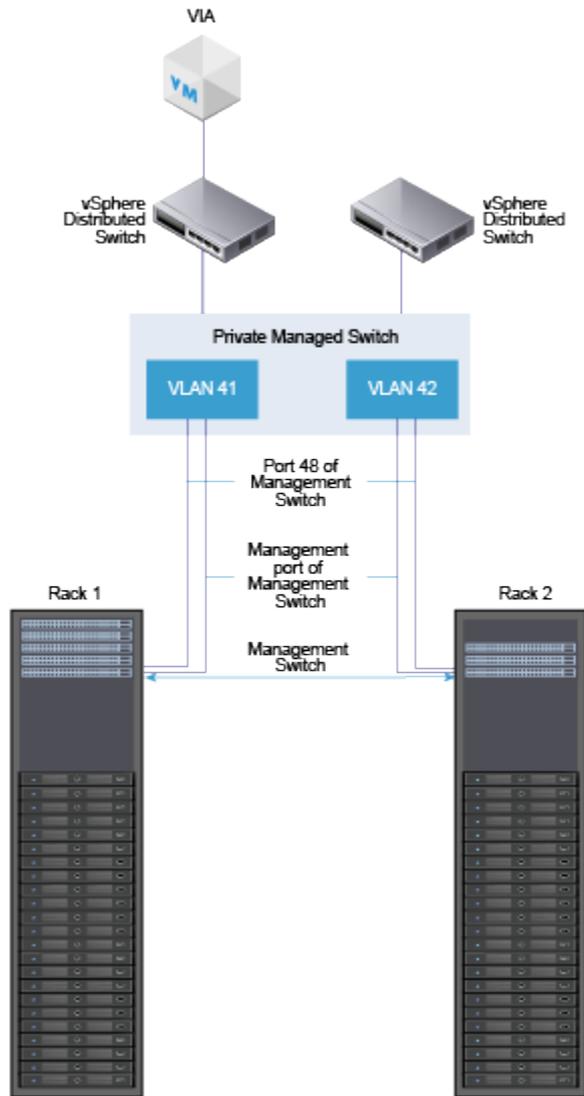
Management Host

If you are using a management host to image the rack, you need a standalone VMware vSphere ESXi 6.0 or later server to host the Windows jump VM. The management host must have at least two NICs, with one NIC connected to the corporate network and one NIC connected to the private network. You also need a 24-port private managed switch.

Table 3-4. VLAN Configuration of the Private Managed Switch

Port	Access Ports
1,2,3,4	VLAN 2000
5,6,7,8	VLAN 2001
9,10,11,12	VLAN 2002
13,14,15,16	VLAN 2003
17,18,19,20	VLAN 2004
21,22,23,24	VLAN 2005

Figure 3-11. Management Host Connection



Private Managed Switch

If this is a multi-rack scenario and the private switch is being shared between racks, configure a private VLAN. For example, create two VLANs in a dual rack environment - VLAN 101 and VLAN 102. VLAN 101 is for rack 1 and VLAN 102 is for rack 2. Port 48 and the management port from the imaging management switch in rack 1 are connected to ports 2 and 3 on the private switch which is configured for VLAN 101. Port 48 and the management port from the imaging management switch in rack 2 are connected to ports 4 and 5 on the private switch which is configured for VLAN 102. The imaging management host is connected to Port 1 which is configured for both VLAN 101 and VLAN 102.

A print out of the VLAN configuration on the imaging management switch should look like this:

```
interface Vlan 1
!untagged GigabitEthernet 0/0-1,6-47
!untagged TenGigabitEthernet 0/48-49
!untagged Port-channel 1-2
!
interface Vlan 2001
no ip address
tagged TenGigabitEthernet 0/48-49
untagged GigabitEthernet 0/2-3
no shutdown
!
interface Vlan 2002
no ip address
tagged TenGigabitEthernet 0/48-49
untagged GigabitEthernet 0/4-5
no shutdown
```

Management Host Settings

Configure the following settings on the imaging management host:

- Install ESXi on the local disk. For the supported version, see the *VMware Cloud Foundation Release Notes*.
- Enable the **Allow virtual machines to start and stop automatically with the system** option.
- Assign the IP address 10.1.0.200 to the vmk0 management network.
- Set the NTP server to the IP address or FQDN of an NTP server that can be accessed from the deployment and can provide NTP service to Linux VMs.

It is important to ensure that the time on the management host is set correctly.

- Enable SSH on the management host.

In a multi-rack scenario, configure an additional vSphere Standard Switch (vDS) for each additional rack. In a dual rack scenario, vSwitch1 should use vmnic1 and should be configured with two Virtual Machine Port Groups (VIA1 and VIA2). The VIA1 port group should be tagged to use VLAN101, and the VIA2 port group should be tagged to use VLAN102. vmnic1 should be connected to the private switch on a port with VLAN101 and VLAN102 visible.

Virtual Machines

The VIA VM runs on the laptop or management host. A jump VM runs on the management host.

If you have multiple physical racks in your environment, you have the following options:

- Image the racks sequentially - image rack 1 first followed by the remaining racks one at a time.
- Image the racks in parallel by configuring a VIA VM per physical rack.

Hardware Configuration

Table 3-5. Jump VM Hardware Configuration

Virtual Hardware	Value
Memory	4 Gb
vCPU	4
HDD	120 Gb
Video card	1 display
SSCI Controller 0	LSI Logic SAS bus sharing: none
Hard disk	120 Gb, Thin Provision
CD/DVD	Client device
Floppy drive	Removed
Network adapters	2 VMXNET3 vNICs
Operating system	Microsoft Windows 7 64-bit or Win2K12
Virtual Machine version	Hardware version 8 or above
Navigate to Options > Advanced > General	Disable logging keyboard.typematicMinDelay = "2000000"

Software Configuration

Perform the following tasks to prepare the jump VM.

- Install a Windows operating system on the VM. This is the version tested with VIA. Other versions may work as well.
- Install VMware Tools.
- Install the latest Windows patches.
- Enables Windows update using the VMware OS Optimization Tool.
- Install the following applications:
 - Firefox or Chrome web browsers
 - PuTTY
 - WinSCP
 - vSphere Web Client (optional)
 - Java Runtime Environment
- If internet access is not available from the Access Virtual Machine, download the executables and binaries for the applications on the VM.
- Verify that Remote Desktop Connection is enabled on the Access Virtual Machine.

Pre-Imaging Checklist

4

You must complete this checklist before beginning the imaging process. It is important that each item in the checklist is set to the specified value, otherwise imaging may fail. You may want to print this checklist and checkmark each row as you verify the setting.

For more information on any item in this table, see [Chapter 3 Setting up your Environment for Imaging](#).

Table 4-1. Pre-Imaging Checklist

Setting	Verified
Review the Bill of Materials (BOM) and verify that there are no discrepancies between the BOM and the hardware being used. If there is a discrepancy, contact VMware Support.	
Review the VMware Compatibility Guide and verify you have the supported server models, disk and storage adapter models, and switch models.	
Verify that the physical racks are wired according to the wiremap. See Rack Wiring .	
Confirm that the inter-rack switches are not connected to the ToR switches. The VIA VM must be connected only to the private network - there should be no external network connections in place before imaging.	
Validate that BIOS Settings for all components are correct. See BIOS Settings .	
Ensure that all servers are powered off. Note This is a new requirement for Cloud Foundation 2.3. All servers must be powered off, otherwise imaging will fail.	
Ensure that switch versions are supported. See Network Switches .	
Verify that firmware settings are set correctly as per the BoM.	
Connect each device in the rack to both PDUs.	
Verify that the cable bend radius is proportionate to the external diameter. See Network Cables .	
Verify that cables are properly routed and labeled.	
Test cables to ensure that installed cabling links provide the transmission capability to support the required data communication.	
Verify that each server has redundant power supplies and that each power supply is connected to a separate rack PDU.	
Ensure that servers and switches have the same airflow direction.	

Table 4-1. Pre-Imaging Checklist (continued)

Setting	Verified
<p>Verify that switches have redundant power supplies and each power supply is connected to a separate power strip.</p>	
<p>If you are using a management host to image your system, ensure that:</p> <ul style="list-style-type: none"> ■ A supported version of ESXi is installed on the management host. For the supported version, see the <i>VMware Cloud Foundation Release Notes</i>. ■ The Allow virtual machines to start and stop automatically with the system option is enabled. ■ IP address 10.1.0.200 is assigned to the vmk0 kernel interface. ■ SSH is enabled on the management host. ■ The access VM, VIA VM, and jump VM meet the required hardware configuration. See Virtual Machines. ■ The required software has been installed on the VMs. See Virtual Machines. 	
<p>Ensure that there are no DHCP or TFTP servers on the network that VIA will connect to for imaging the rack. VIA uses DHCP to assign IP addresses to the rack components, and any external DHCP servers can cause imaging to fail. Disconnect the uplink cables on the TORs to limit exposure.</p>	

Installing VIA

5

You can install VIA on a desktop, laptop, or an ESXi host, also referred to as the management host. A laptop or desktop is recommended when you are imaging a single rack. A management host is better suited for an environment where you have several physical racks in your datacenter.

This chapter includes the following topics:

- [Installing VIA on a Laptop or Desktop](#)
- [Installing VIA on a Management Host](#)

Installing VIA on a Laptop or Desktop

VIA is a virtual appliance. After you install the VIA VM on your laptop, you copy the software bundle to the VIA VM. You can then access the VIA user interface through a browser on the laptop.

Prerequisites

- Ensure that you have the infrastructure for VIA available and that you have set up your physical environment as described in [Chapter 3 Setting up your Environment for Imaging](#).
- Download the VIA OVF file, Cloud Foundation software bundle, and the md5sum file on the laptop or desktop.

Procedure

- 1 Connect one port of the network adapter to your laptop and one port to the unmanaged switch.
- 2 Connect two ports of the unmanaged switch as follows:
 - one port to the ethernet management port of the management switch
 - one port to port 48 of the management switch
- 3 Connect the .iso image to the VM's CDROM drive.
Follow the wizard prompts to specify where to save the OVF file and accept the license agreement.
- 4 Configure time settings on the laptop.

- 5 Upload the Cloud Foundation software bundle on to the VIA VM by pointing the CD /DVD drive to the bundle ISO. Ensure that the CD/DVD drive is connected.
- 6 Configure network settings on the laptop.
 - a Connect one NIC on the laptop to the corporate network and the other to the unmanaged switch.
 - b Manually assign IP address 192.168.100.190 to the interface facing the unmanaged switch on the laptop.

This allows for separation of network traffic between the corporate network and the private network that is established between the physical rack and VIA. It also helps ensure that the DHCP service which is part of is VIA is confined to the private network between the physical rack and VIA.

- 7 For the browser on the laptop that will be used to access VIA, make the following selections.
 - In Network Connection, disable the proxy.
 - Select **Auto-detect proxy settings for this network** so that the browser detects the proxy settings for your network.
- 8 Power on the VIA VM.
- 9 Ensure that you can ping the VIA VM (IP address is 192.168.100.2) from the laptop.

What to do next

Open a web browser on the laptop and type the following URL to connect to VIA:

`http://192.168.100.2:8080/via/`

Installing VIA on a Management Host

Prerequisites

- Ensure that you have the infrastructure for VIA available and that you have set up your physical environment as described in [Chapter 3 Setting up your Environment for Imaging](#).
- Download the VIA OVF file, Cloud Foundation software bundle, and the md5sum file on your local file system.

Procedure

- 1 Deploy the VIA OVF file on the management host.
 - a Login to the vSphere Web Client on the management host.
 - b Right-click the management host and click **Deploy OVF Template**.
 - c In source location, select **Local file**. Click **Browse** and select the VIA OVF from your local file system.
 - d Click **Next**.

- e Review the OVF file details and click **Next**.
 - f Accept the OVF license agreements and click **Next**.
 - g Specify a name and location for the OVF and click **Next**.
 - h Select a resource and click **Next**.
 - i Select the disk format to store the VIA disks and the datastore to store the deployed OVF template and click **Next**.
 - j On the Setup networks page, connect VIA to the private switch connected to rack 1.
 - k Review the deployment details and click **Finish**.
- 2** Copy the Cloud Foundation bundle to the management host.
- a On the management host, create a single datastore named datastore1.
 - b In datastore1, create a folder named ISO bundle and copy the Cloud Foundation bundle file to this folder.
- 3** Configure time settings on the management host.
- a In the vSphere Web Client, navigate to the management host in the vSphere inventory.
 - b Select **Manage** and then select **Settings**.
 - c Under **System**, select **Time configuration** and click **Edit**.
 - d Select **Manually configure the date and time on this host**.
 - e Set the time and date manually.
 - f Click **OK**.
- 4** Upload the software bundle on to the VIA VM.
- a Right-click the VIA VM and select **Edit Settings**.
 - b Click the **Hardware** tab and select the CD/DVD drive.
 - c Select the **Connected** check box to connect the CD.
 - d Select **Connect at power on** so that the CD-ROM drive is connected when the virtual machine starts.
 - e Select **Datastore ISO** under **Device Type**.
 - f Click **Select**, browse to the ISO Bundle folder in datastore1 on the management host, and select the bundle.
 - g Click **OK**.
- 5** Create a VM on the management host to serve as the jump VM.
- Connect one NIC on the jump VM to the network and the other to the private managed switch.

The jump VM must have a static IP address. The IP range 192.168.100.151 to 192.168.100.199 is usually available for the jump VM.

This allows for separation of network traffic between the datacenter network and the private network that is established between the physical rack and VIA. It also helps ensure that the DHCP service which is part of is VIA is confined to the private network between the physical rack and VIA.

- 6 Copy the md5sum file on the jump VM.
- 7 For the browser on the jump VM that will be used to access VIA, make the following selections.
 - In Network Connection, disable the proxy.
 - Select **Auto-detect proxy settings for this network** so that the browser detects the proxy settings for your network.
- 8 Ensure that you can ping the VIA VM (IP address is 192.168.100.2) from the jump VM.
If you cannot ping the VIA VM, check the route on the jump VM.
- 9 Power on the VIA VM.

What to do next

Open a web browser and type the following URL to connect to VIA:

<http://192.168.100.2:8080/via/>

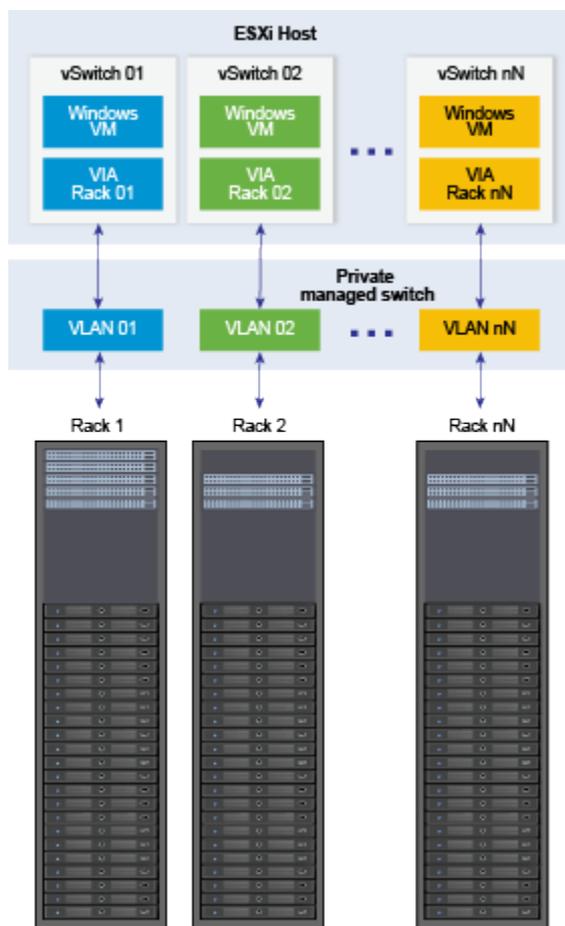
Imaging Physical Racks

6

When you image a physical rack, the software in the Cloud Foundation bundle is loaded onto the physical rack.

In a multi-rack environment, you can either image all racks in parallel, or image rack 1 first followed by the other racks one at a time. To image multiple racks in parallel, you need a vSphere Distributed Switch and VIA VM for each rack.

Figure 6-1. VIA Setup for Parallel Imaging of Multiple Physical Racks



This chapter includes the following topics:

- [Image a Physical Rack](#)

- [Resume Imaging](#)
- [Image Additional Racks](#)

Image a Physical Rack

VIA images the rack components in a pre-determined order, which is determined by the availability of network route to the different components of the rack.

All switches are imaged first. This enables VIA to access the servers through the switches for imaging. The imaging order is as follows.

1 Management switch

The management switch is the main access gateway through which the Cloud Foundation management data is routed. The management ports of the ToR and inter-rack switches are connected to the management switch. The data ports of the ToR switches are also connected to the management switch. VIA is connected to the rack through a designated port on the management switch. It is therefore required that the management switch is the first component imaged by VIA in order to obtain access to the other components of the rack.

2 ToR switches and Inter-rack switches (inter-rack switches are on rack 2 only)

ToR and inter-rack switches are imaged in parallel.

ToR switches provide connectivity to servers in each rack out to inter-rack switches. The first pair of ToR switches provide connectivity to your datacenter network. Inter-rack switches inter-connect multiple racks enabling a scale out architecture for the datacenter. They create an stretched L2 backplane between racks.

3 Servers

Servers must be rebooted after the switches are imaged. They are then imaged in parallel.

For each component that is being imaged, the following tasks are performed.

1 Discovery

Rack components are discovered using the DHCP service running on the VIA VM. The DHCP Service uses the device type information to identify the device being discovered. Apart from the device type information, the DHCP service also uses hardware vendor specific strings to determine whether the switch being imaged is a management, ToR, or inter-rack switch.

The first component to be discovered is the management switch. The DHCP service hands out a pre-determined IP address for the management switch followed by a PXE image specific to the management switch.

After the management switch is imaged, the ToR and inter-rack switches are discovered and imaged. The ToR switch enables discovery of the data network of the servers which is used to receive the installation image delivered by the DHCP service.

2 Image installation

Image installation refers to installing software on the components to make them operational. The software depends on the component type - an Operating System for switches and a Hypervisor for servers.

3 Configuration

This step in the imaging process ensures that the components of the rack work like a homogenous system. Configuration of each rack component is different. If any configuration step fails for the management, ToR, or inter-rack switches, the imaging run stops at that point. You must resolve the issue before you can proceed with imaging. If a configuration step fails for a server, you can skip that server and continue imaging the remaining servers.

Procedure

1 Upload Software Bundle

The software bundle ISO file contains the software bits and scripts to be imaged on the physical rack. You can upload multiple bundles at a time and activate the bundle that is to be used for imaging.

2 Add Server VIBs

If you have servers in your datacenter that need VIBs not included in the core bundle, you can upload these additional VIBs on VIA.

3 Add Custom ESXi Installer ISO

You can choose to image the servers in your rack with an ISO different from the one included in the core software bundle. The ISO must be of the same version that is included with the Cloud Foundation software bundle. For more information, see the Cloud Foundation Release Notes.

4 Specify Imaging Details

At the Details step of an imaging run, you provide a name and description for the imaging run as well component and port information for the rack.

5 Monitor Imaging

In the Monitor Imaging step of the imaging workflow, you can see the imaging status on all devices in your physical rack. You must reboot the servers after the switches are imaged.

6 Verify Inventory

In the Verify step of the imaging workflow, the system collects inventory information for each device in the rack.

7 Post Imaging Checks

In the final step of the imaging workflow, VIA creates a rack inventory file.

8 Download Inventory File

You must download and save the inventory file for each imaged rack or host. This file is required during rack bring-up and while adding a host. For the first rack, the inventory file is copied over to the SDDC Manager VM during imaging. However, you must still download the file and save it to an accessible location in case the file on is corrupted for some reason. For additional racks, you must upload this file manually while adding the rack to your system.

Upload Software Bundle

The software bundle ISO file contains the software bits and scripts to be imaged on the physical rack. You can upload multiple bundles at a time and activate the bundle that is to be used for imaging.

The bundle contains the following software:

- SDDC Manager
- vSphere (Platform Services Controller, vCenter Server and ESXi)
- vSAN
- NSX
- vRealize Suite (vRealize Log Insight, vRealize Operations, and vRealize Automation)
- VMware Horizon
- App Volumes

You can upload additional VIBs based on server models and can also provide an ESXi ISO different from the one included in the core bundle to image the servers in your racks.

Prerequisites

- Download the Cloud Foundation software bundle and the md5sum file on your laptop, desktop, or jump VM.
- If you are re-purposing hosts in your datacenter, backup the data on the hosts. They are wiped clean during imaging.
- Ensure that the KVM console for servers is closed.

Procedure

- 1 In a browser window on the jump VM, type `http://192.168.100.2:8080/via`.

- Click **Bundle**. Ensure that you are in the Core Bundle tab.

- In the **Bundle Location** area, click **Refresh**.

Wait for the message **CD mounted successfully** to be displayed.

- In the **Bundle Hash** area, click **Browse**, navigate to the directory that contains the MD5SUM file, select the file, and click **Open**.
- Click **Upload Bundle**.

The bundle upload can take several minutes. After the upload is complete, the message **Bundle uploaded successfully** is displayed in the Upload Bundle area.

- In the **Bundle Info** area, select the bundle in **Available Versions** and click **Activate Bundle**.

The selected bundle is now the active bundle for imaging and is ready to be used. Active bundle details are displayed next to **Active Bundle**.

The bundle is copied to the `/mnt/cdrom/` directory on the VIA VM.

Add Server VIBs

If you have servers in your datacenter that need VIBs not included in the core bundle, you can upload these additional VIBs on VIA.

Prerequisites

Download the 3rd party VIB you want to add to the software bundle on the jump VM where you are running VIA.

Procedure

- 1 In the Modify VIBs tab, select the vendor or click **Add Vendor**.

- 2 For a new vendor, type the vendor name and click **Add**.
- 3 Select the server model or click **Add Model**.
- 4 For a new model, type the model name.

The model name you enter must match the model or product name on the server's BIOS settings. This ensures that the correct VIB is loaded on the server.

The available VIBs are displayed.

- 5 To add 3rd party VIBs, click **Browse** next to the **Select VIB** field, select the VIB and click **Upload VIB**.

Note Only files with a .vib extension are accepted.

- 6 In **Available VIBs**, select the VIB to be added to the bundle.
- 7 Click **Upload VIB**.

Add Custom ESXi Installer ISO

You can choose to image the servers in your rack with an ISO different from the one included in the core software bundle. The ISO must be of the same version that is included with the Cloud Foundation software bundle. For more information, see the Cloud Foundation Release Notes.

Prerequisites

Download the ISO you want to use for imaging on the jump VM where you are running VIA.

Procedure

- 1 In the Bundle tab, click **Modify ISOs**.
 - ◆ Active

Vendor Name :

Available ISOs :

Name	Source	Active
ESXi-6.5.0-20170702001-vxrack_d14g_p01-standard.iso	Bundle	<input checked="" type="radio"/>

Select ISO to Add :

MD5 Checksum :

Checksum Type : MD5 SHA-1

2 Select the vendor name.

The ISO in the core bundle is displayed.

3 In **Select ISO to Add**, click Browse, select the ISO to add and click **OK**.

4 Type the MD5 or SHA-1 checksum for the ISO.

You can either download the MD5 checksum from the location where you downloaded the ISO, or generate it via another tool.

5 Click **Upload ISO**.

6 In the **Available ISOs** section, click the **Active** radio button next the newly uploaded ISO.

Specify Imaging Details

At the Details step of an imaging run, you provide a name and description for the imaging run as well component and port information for the rack.

Prerequisites

- Software bundle must have been uploaded and activated.
- All servers are powered off.
- You must be able to power on the servers one at a time.
- All items in the pre-imaging checklist have been completed.

Procedure

- 1 In the VIA user interface, click **Imaging**.

VIA Bundle **Imaging** Inventory History Logs About  (0)

Details **Imaging** Verify Finish

Name

Description

Deployment Type

Rack Type

MANAGEMENT SWITCH 

1 Vendor: Model: IP: MAC:

TOR SWITCH 

1 Vendor: Model: IP: MAC:

2 Vendor: Model: IP: MAC:

INTER-RACK SWITCH Number: 

ESXI SERVER Number: 

1 Vendor: Model: IP: MAC:

2 Vendor: Model: IP: MAC:

3 Vendor: Model: IP: MAC:

Primary ESXi Server Info

If you wish to specify which server becomes the primary ESXi server, please select one of the data types below and enter the corresponding value for your preferred server. This step is optional: if no value is given here then the first server to boot is chosen to be the primary ESXi server.

Connected TOR Port Number MAC Address

TOR Port:

Start Imaging

- 2 (Optional) Type a name and description for the imaging run.

It is recommended that you add the rack serial number or other rack identification details in the Name or Description field. The name and description is recorded in the imaging details JSON file after imaging is complete, which helps identify the rack to which the imaging details file belongs.

- 3 In **Deployment Type**, select **Cloud Foundation Full Deployment**.
- 4 In **Rack Type**, select **First Rack** if you are imaging rack 1. For additional racks, select **Additional Rack**.

- 5 For the management switch and ToR switches, select the vendor and model.
The IP address for each switch is displayed.
- 6 (Optional) Type the MAC address of each switch.
- 7 If the physical rack contains inter-rack switches (rack 2 only), type the number of inter-rack switches in the **Number** field in the **Inter-rack Switch** section.
- 8 Select the vendor and model number of the inter-rack switches.
- 9 In the **Server** section, type the number of servers in the physical rack.
- 10 Select the vendor for each server.
- 11 You can specify the server where the SDDC Manager VMs will be deployed. This server is referred to as the primary host. Either specify the port number on the ToR switch that the server is connected to, or the MAC address of the server.

If you do not type the port number or MAC address, the SDDC Manager VMs are deployed on the first server to be rebooted.
- 12 Click  in any section to view the VIA properties file.

The VIA properties file displays rack specification values from the activated software bundle. If required, edit the properties as appropriate. For example, delete the ToR ports that you are not using to speed up the imaging process.
- 13 Click **Save**.
- 14 Click **Start Imaging**. The wiring diagram for the rack is displayed.

Review the wiring diagram to check if your rack is wired according to the displayed wiremap. See [Rack Wiring](#).

Click **Confirm** if it is accurate. If you need to make any wiring changes, click **Cancel**.

What to do next

Once imaging starts, notifications (errors, warnings, and information) are displayed in the top right corner of the VIA window. Click the  icon to review the messages so that you can take the appropriate steps to complete the imaging successfully.

Monitor Imaging

In the Monitor Imaging step of the imaging workflow, you can see the imaging status on all devices in your physical rack. You must reboot the servers after the switches are imaged.

Procedure

- 1 Click the **Imaging > Imaging** tab.



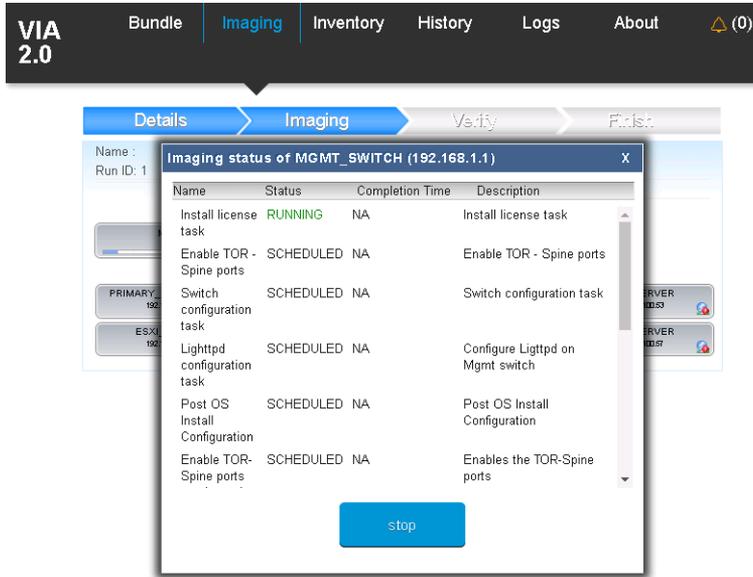
The run details, rack details, and imaging status for switches and servers in the rack are displayed. The devices in the physical rack are displayed in the order in which they will be imaged. IP addresses for each device are displayed as well.

For information on next steps if a device fails to be imaged, see [Resume Imaging](#).

- 2 Note the **Run ID**.
- 3 After the switches are imaged, the UI displays a notification to reboot the servers. Reboot the servers within 2 hours. If the session times out before the reboot is completed, click **Resume** to continue imaging.

Once server imaging begins, the IP address for the primary ESXi host is displayed. Note down this address since you may need to log in to this host.

- 4 Click a device to see information about the imaging tasks completed and in-progress tasks.



It can take approximately 95 minutes for rack 1 to be imaged. During imaging, the password of all rack components except SDDC Manager is set to `EvoSddc!2016`. The SDDC Manager password is set to `vmware1234`. Note that the ToR and inter-rack switches are named as follows:

- ToR switch connected to port 41 of the management switch is named TOR20
- ToR switch connected to port 42 of the management switch is named TOR21
- Inter-rack switch connected to port 43 of the management switch is named INTER-RACK30
- Inter-rack switch connected to port 44 of the management switch is named INTER-RACK31

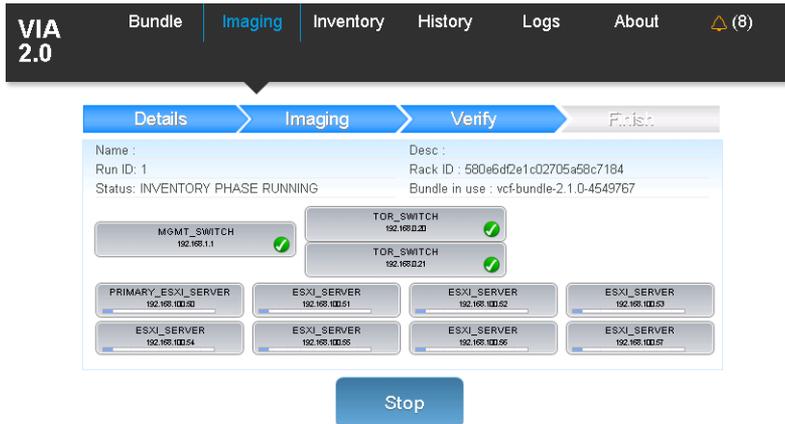
For information on next steps if a device fails to be imaged, see [Resume Imaging](#).

Verify Inventory

In the Verify step of the imaging workflow, the system collects inventory information for each device in the rack.

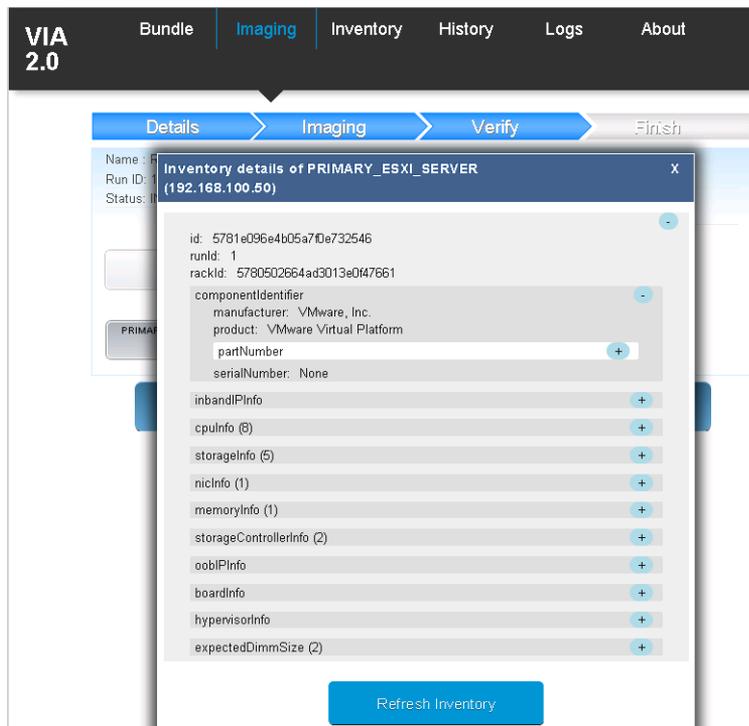
Procedure

- 1 Click the **Imaging > Verify**



The status of inventory collection on each device in the rack is displayed.

- 2 Click each component to see its inventory information. You can expand a component to see



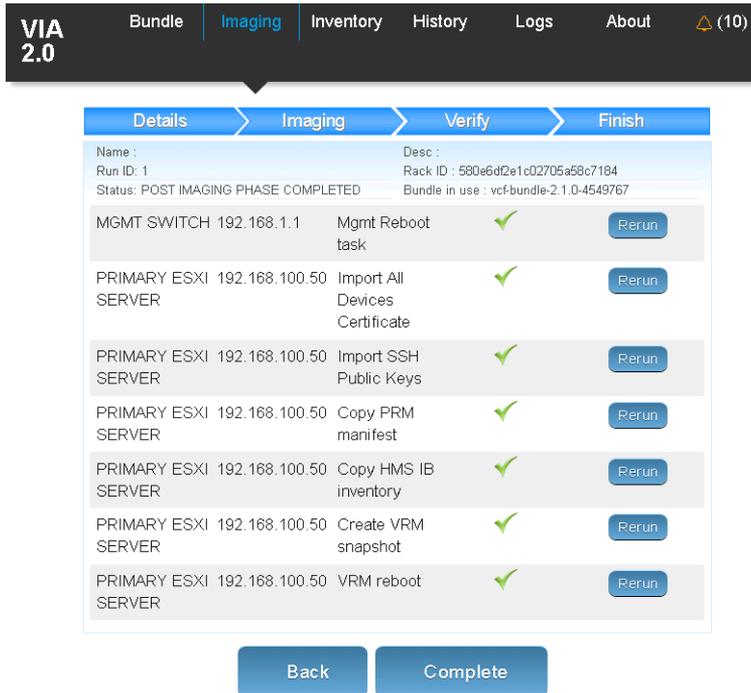
more details.

Post Imaging Checks

In the final step of the imaging workflow, VIA creates a rack inventory file.

Procedure

- 1 Click the **Imaging > Finish** tab.



Post imaging tasks are displayed.

- 2 If a task is not completed successfully, click **Rerun**.
- 3 After each displayed task has an ✓ icon next to it, click **Complete**.

The rack inventory file is created for the customer. This file includes the SDDC Manager password generated during imaging. The imaged rack is now ready to be shipped to the customer.

- 4 Power down the primary rack.

It is important to power down the rack even if you are deploying Cloud Foundation on a ready system so that the management switch is rebooted.

Download Inventory File

You must download and save the inventory file for each imaged rack or host. This file is required during rack bring-up and while adding a host. For the first rack, the inventory file is copied over to the SDDC Manager VM during imaging. However, you must still download the file and save it to an accessible location in case the file on is corrupted for some reason. For additional racks, you must upload this file manually while adding the rack to your system.

Procedure

- 1 Click the **Inventory** tab.
- 2 Select the Run ID.

3 Click **Download**.

The file is download on the jump VM. Each inventory file is identified by the ImagingID. The name and description in the file match the details you specified in the Imaging tab.

4 Copy the downloaded file to an accessible location.

5 Rename the file such that the file name indicates the rack or host that the inventory file belongs to.

Resume Imaging

A device may fail to be imaged because of possible hardware faults or mis-configuration, or network issues. You can take a number of actions that can help in continuing with the imaging run.

Fix Issues During the Monitorf you are able to resolve the hardware problem, click r Imaging Step

During the monitor step in the imaging workflow, you can identify imaging failures by looking at the progress bar on the components in the **Imaging > Imaging** tab. An  icon indicates that it has been imaged successfully. An  icon indicates that one or more imaging tasks on that devices failed.



1 Click the component to display the imaging task list for that device. Then do one of the following:

- Click **Retry** to re-start imaging on that device .
- Click **Remove** to remove that device from the VIA UI and database and then click **Yes** to confirm. The removed device is grayed out and it is not imaged. Ensure that you remove this device from the physical rack before shipping it to the customer. To add a removed device back to the VIA UI, click the device and click **Add to Inventory**. The device is added back to the VIA UI and database.

If the primary ESXi server fails to be imaged, you cannot resume imaging by removing it. It is mandatory for the primary server to be imaged.

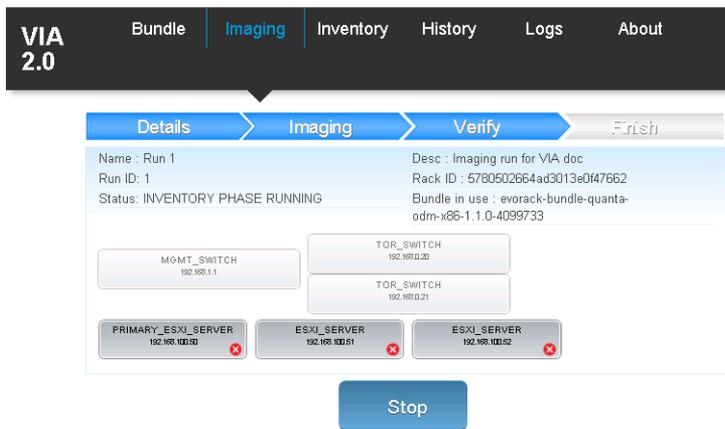
- 2 If you need to resolve a hardware issue before re-trying imaging on that device, close the task list dialog box. In the **Imaging > Imaging** window, click **Stop**.
- 3 If you are able to resolve the hardware problem, click **Resume**. Imaging is resumed from the state where it had stopped. If you need additional time to resolve the hardware issue or there are other hardware problems, click **Abort**. The imaging run is discarded.

If you are unable to resolve the hardware issues, contact VMware Support.

- 4 Click **Next** to proceed to the next step in the workflow.

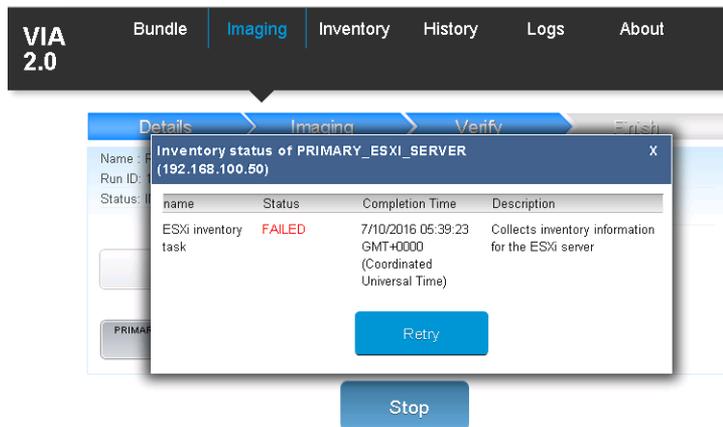
Fix Issues During the Verify Imaging Step

During the verify step in the imaging workflow, you can identify imaging failures by looking at the progress bar on the components in the **Imaging > Verify** tab. An  icon indicates that inventory information has been collected successfully. An  icon indicates that the tasks on that device



failed.

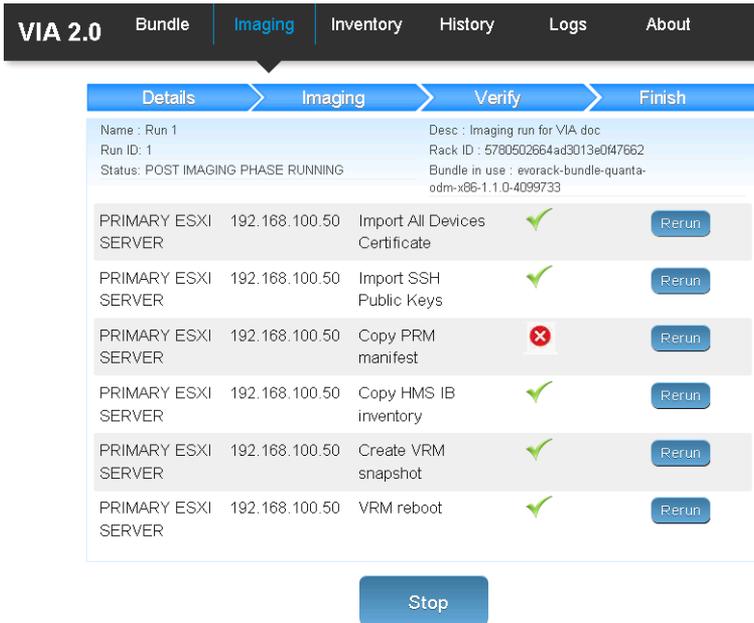
- 1 Click the component to display the verification task list for that device.



- 2 Click **Retry**
- 3 Once the device displays an  icon, click **Next**.

Fix Issues During the Finish Imaging Step

During the finish step in the imaging workflow, failed post-imaging tasks are displayed with an .



VIA 2.0 Bundle **Imaging** Inventory History Logs About

Details > **Imaging** > Verify > Finish

Name : Run 1 Desc : Imaging run for VIA doc
 Run ID: 1 Rack ID : 5780502664ad3013e0f47662
 Status: POST IMAGING PHASE RUNNING Bundle in use : evorack-bundle-quanta-odm-x86-1.1.0-4099733

PRIMARY ESXI SERVER	192.168.100.50	Import All Devices Certificate	✓	Rerun
PRIMARY ESXI SERVER	192.168.100.50	Import SSH Public Keys	✓	Rerun
PRIMARY ESXI SERVER	192.168.100.50	Copy PRM manifest	✗	Rerun
PRIMARY ESXI SERVER	192.168.100.50	Copy HMS IB inventory	✓	Rerun
PRIMARY ESXI SERVER	192.168.100.50	Create VRM snapshot	✓	Rerun
PRIMARY ESXI SERVER	192.168.100.50	VRM reboot	✓	Rerun

Stop

icon.

- 1 Click **Rerun** to run the failed task again.
- 2 After all tasks display an  icon, click **Complete**.

Opening a Cancelled Run

If you had accidentally cancelled an imaging run, you can re-open it.

- 1 In the VIA user interface, click **History**.
- 2 In the **Select Run ID** drop-down, select the run ID you want to open.
- 3 Click **Reopen**.

The selected run is opened in the state it was at the time the run had been cancelled.

Image Additional Racks

Follow this procedure for each additional rack if you are imaging racks incrementally in a multi-rack environment.

Procedure

- 1 Disconnect port 48 of the management switch on rack1 from the private managed switch.
- 2 Connect port 48 of the management switch on the next rack to the private managed switch.
- 3 Follow [Image a Physical Rack](#).

Imaging Individual Devices

7

You can image a server or management switch as an individual device.

This chapter includes the following topics:

- [Image Individual Server](#)
- [Image New Management Switch](#)

Image Individual Server

If a server fails to be imaged, you can image that server as an individual device rather than re-imaging the complete rack. Or you can image a new or replacement server before adding it to a rack.

Prerequisites

- Mount the host in the appropriate slot in the physical rack. For a replacement host, mount it in the same slot as the previous host and wire it according to the same wiring connections.
- If VIA is installed on a laptop, connect the NIC port on the laptop to port 48 of the management switch on which the host is being imaged. If VIA is installed on a management host, connect the host to a private managed switch that is connected to port 48 of the management switch on which the host is being imaged.
- VIA must have access to the inband network only.
- Software bundle must have been uploaded. See [Upload Software Bundle](#).
- BIOS settings must have been set on the server to be imaged. See [BIOS Settings](#).
- Server must be in PXE boot mode.
- Power off the server before imaging.

Procedure

- 1 In the VIA user interface, click **Imaging**.
Ensure that you are in the **Details** tab.
- 2 (Optional) Type a name and description for the imaging run.
- 3 In **Deployment Type**, select **Cloud Foundation Individual Deployment**.

- 4 In **Device Type**, select **ESXi_SERVER**.
- 5 In **Rack Type**, select **Primary Rack** if you are imaging a server in rack 1. For a server in an additional racks, select **Add-On Rack**.
- 6 Select the vendor and model number of the server.
The IP address of the server is displayed.
- 7 If you are imaging a host for a multi-rack Cloud Foundation deployment, you must type the MAC address of the host being imaged. This ensures that VIA images the correct host.
- 8 Click **Start Imaging**.
- 9 Change the boot device for the next boot to **PXE Mode**.

Note The steps below document the procedure for Dell servers. For servers from other vendors, refer to the vendor documentation.

- a In a web browser, open the Integrated Dell Remote Access Controller page and login with default credentials (root/calvin).
 - b In the Properties tab, click **Launch** in the Virtual Console Preview.
The console opens in a new window.
 - c Open the keyboard.
 - d In the Properties tab, click **Power Cycle System** (cold boot) in the Quick Launch Tasks
 - e In the console window, press F11 to access the Boot Manager.
F11 = Boot Manager is highlighted.
 - f Select One-shot BIOS Boot Method
 - g Select PXE Boot.
Do not power cycle or reset the server.
You can disable the PXE mode after imaging is complete.
- 10 Power on the server.
The server is discovered and the imaging process begins.
 - 11 Open the KVM console to the server.
 - 12 Monitor the ESXi installation on the console.
The server is rebooted after ESXi installation is complete. Ensure that ESXi is booting from the installed copy of ESXi and not from the network.
After the server boots from the installed copy of ESXi, VIA continues imaging the server.

13 Download the inventory file.

You need this file when adding an imaged host to a Cloud Foundation rack.

- a Click the **Inventory** tab.
- b Select the Run ID.
- c Click **Download**.

The file is download on the jump VM.

- d Copy the downloaded file to an accessible location.

14 Disconnect VIA and shutdown the laptop or management host.

Image New Management Switch

Imaging the new management switch with VIA installs the necessary software on the switch.

Prerequisites

- Management switch must be connected to the laptop or management host where VIA is installed.
 - If VIA is installed on a laptop, the NIC port on the laptop must be connected to port 48 of the management switch.
 - If VIA is installed on a management host, the management host must be connected to a private managed switch that is connected to port 48 of the management switch.
- Identify the Cloud Foundation version in your environment and ensure that the appropriate bundle and md5sum file is uploaded on VIA.

Note Do not connect the management switch to any host before or during imaging.

Procedure

- 1** In the VIA user interface, click **Imaging**.
Ensure that you are in the **Details** tab.
- 2** (Optional) Type a name and description for the imaging run.
- 3** In **Deployment Type**, select **Cloud Foundation Individual Deployment**.
- 4** In **Device Type**, select **MGMT_SWITCH**.
- 5** Select the vendor and model number of the switch. The IP address is displayed.
- 6** Click **Start Imaging**.
- 7** Disconnect the switch from the laptop or management host.

Review Alarms and Notifications



After the imaging run is completed, you must review all alarms and notifications by clicking the bell icon at the top right corner of the VIA window.

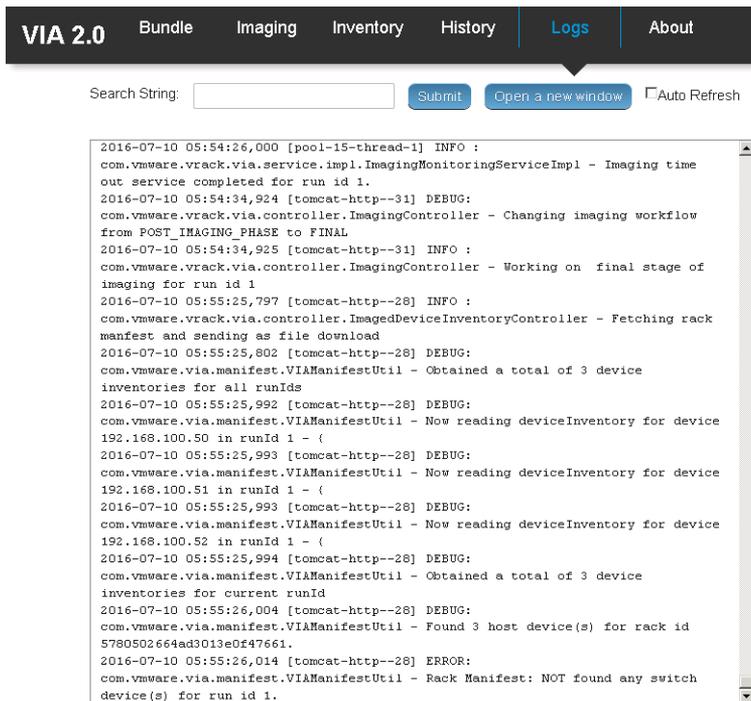
Viewing the VIA Log File

9

The log file displays information for all VIA services.

Procedure

- ◆ On the left navigation bar in the VIA user interface, click **Logs**.



A consolidated log of VIA services is displayed sorted by the time stamp. A maximum of 500 entries is displayed at a time.

You can filter the logs by typing a search string and clicking **Submit**. For example, you can search for activities on the primary ESXi server.

To display the complete log file, click **Open a new window**.

The **Auto Refresh** option is selected by default where the log file automatically scrolls to display the most current information.

Viewing Results of an Imaging Run

10

You can view the imaging history for an imaged rack or the status of individual devices on an imaged rack.

This chapter includes the following topics:

- [View Imaging History](#)
- [View Inventory](#)

View Imaging History

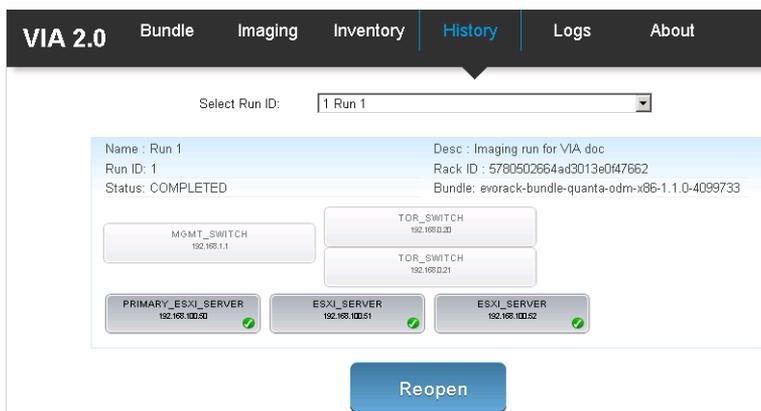
You can view the status of an imaging run by specifying its run ID if the rack state has not changed since that run was completed. If you imaged multiple racks using the same VIA VM, you can view the imaging history of each rack by specifying its run ID.

Prerequisites

Verify that an imaging run is not in progress.

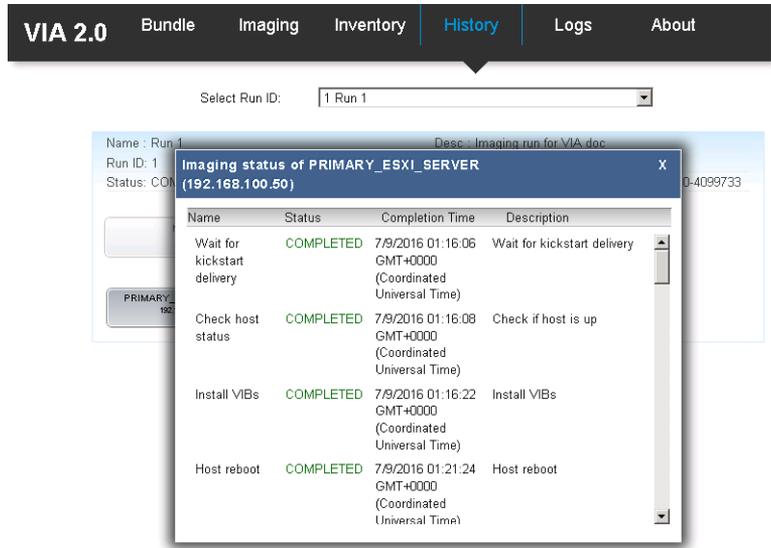
Procedure

- 1 In the VIA user interface, click **History**.



- In **Select Run ID**, select the run ID for which you want to view the imaging history. You can only view the history for a run if the state of the rack has not changed since the run was completed.

Imaging history appears for all devices that are imaged during the specified run.



- To view details for a device, click the expand icon next to the device.
- To reopen a previous run, select the run ID and click **Reopen**.
You can continue imaging a cancelled run by reopening it.

View Inventory

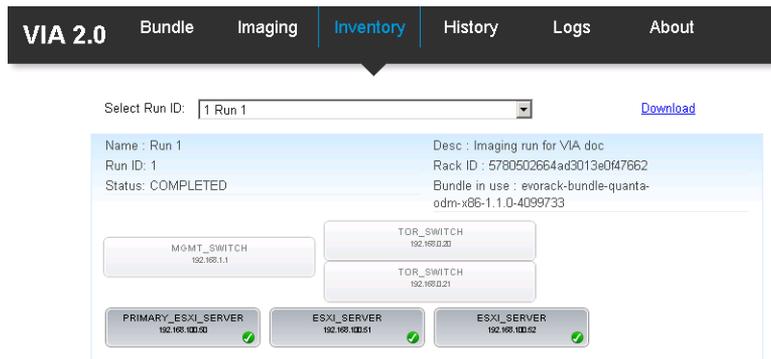
The Inventory page displays a consolidated report of the rack inventory. You can view device details by expanding the appropriate device.

Prerequisites

Verify that an imaging run is not in progress.

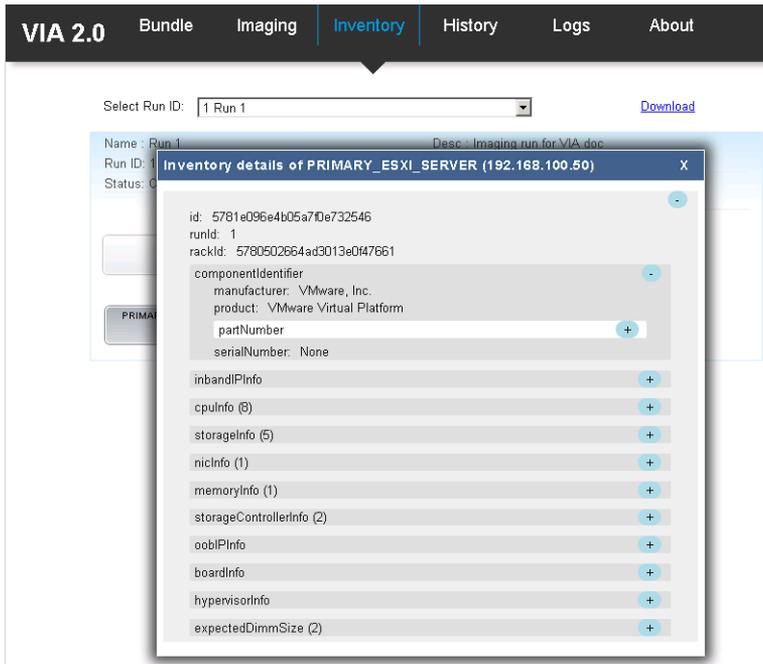
Procedure

- In the VIA user interface, click **Inventory**.



- In Select Run ID, select run ID.

The device inventory for the selected imaging run is displayed.



- To view details for a device, click the expand icon next to the device.
- To download the rack inventory click **Download Manifest** and specify the directory where the file is to be saved.

The device inventory is saved as a JSON file.

Guidance on Server OOB Port Management

11

Cloud Foundation does not manage OOB ports for 2.3 and later releases. You can manage the OOB ports per your needs. You must make a note of IP addresses assigned to servers and OOB passwords.

OOB Ports Wiring to Management Switch

This section describes ways in which you can wire the OOB ports to the management switch.

- Connect server OOB ports to an external OOB management switch. This is the recommended method.
- Connect server OOB ports to the management switch in the rack and assign static IP addresses to them. If you choose to use the 192.168.0.0/22 network, ensure that the IP addresses are within the range of 192.168.1.2 to 192.168.3.253. You cannot use IP address outside this range.

OOB Ports Configuration

After imaging is complete, add a route to your jump VM to connect to the OOB ports on the servers.

Procedure

- ◆ Add a route on the laptop/desktop or jump VM where you have deployed VIA to allow access to the physical servers. Here is an example on how to add a route on a Windows jump VM for a 192.168.0.0 network.

```
route add 192.168.0.0 mask 255.255.255.0 192.168.100.1 if 16
```

where *16* is the ID for rack 1. To find the interface number on a Windows laptop or desktop , follow the steps below.

- a In a command window, type the command **netsh**.
- b Type the command `int ipv4 show interfaces`.

Troubleshooting VIA

12

This chapter includes the following topics:

- [ESXi Server has Incorrect BIOS Settings](#)
- [ESXi Server has Bad SD Card](#)
- [Management Switch Boots into EFI Shell](#)
- [A switch did not start imaging](#)
- [One or More Servers Are Not Discovered](#)
- [Server Imaging Failed at Kickstart Delivery Task](#)
- [Server Imaging Failed at Check Host Status Task](#)

ESXi Server has Incorrect BIOS Settings

Problem

Host failed to be imaged with the message **Post install reboot ESXi task failed.**

Cause

ESXi server has incorrect BIOS settings.

Solution

- 1 Check the ESXi server console.
- 2 If the console displays a gray screen with the message Unable to find boot device, check that the BIOS setting is SATADOM for Quanta servers, and SD card or SATADOM for Dell servers.
- 3 Fix the hardware problem.
- 4 On the **Imaging** tab, click the host that displayed the red icon and click **Retry**.

ESXi Server has Bad SD Card

Problem

Device failed to be imaged with the message **Kickstart image not delivered..**

Cause

ESXi server has bad SD card.

Solution

- 1 Replace the SD flash card in the ESXi server.
- 2 On the **Imaging** tab, click the host that displayed the red icon and click **Retry**.

Management Switch Boots into EFI Shell

Problem

After rebooting, the management switch boots into EFI shell instead of ONIE mode.

Note This issue affects only management switches running the Cumulus OS.

Cause

The switch was not in ONIE mode and after rebooting, it boots into an EFI shell.

Solution

- 1 Connect to the management switch with a console cable.
- 2 Press **DEL** to change the boot order.
- 3 Select **P0**.
- 4 Select **Save changes**.
- 5 Select **Save changes and restart**.
- 6 To wipe the switch login as cumulus, type `sudo cl-img-select -k`.

A switch did not start imaging

Imaging did not begin on a switch.

Problem

Cause

The switch mode was not set correctly prior to imaging.

Solution

- ◆ Using the serial console port, check the switch mode and ensure it matches the mode specified in [Network Switches](#).

One or More Servers Are Not Discovered

One or more servers are not discovered by VIA.

Problem

When servers are not discovered, the corresponding UI objects on the **Imaging > Imaging** tab do not turn green.

Cause

There are several reasons due to which servers might not be discovered:

- Server was not powered on.
- Server was powered on, but is stuck in BIO mode due to a hardware problem.
- Server was not properly connected to the ToR switches.
- Server was powered on, but is not trying to PXEBoot from the appropriate network card for VIA to send the ESXi image.

To identify the servers which are not discovered in VIA, you must search through the VIA log file.

Solution

- 1 In the VIA user interface, click **Logs**.
- 2 Type one of the following search strings. All search strings must be entered without quotes.

- Identifying undiscovered servers by their assigned BMC IP addresses

Type `IPMI BMC IP Address =`

All discovered servers display the corresponding BMC IP address programmed on those servers. The BMC IP of the servers which were not discovered do not appear here. Log in to the KVM console of the undiscovered BMC consoles for further troubleshooting.

- Identifying undiscovered servers based on server port connection to the ToR switches

Type `getToRPortNum`

Displays discovered servers with their ToR port connectivity. Use the process of elimination to identify the servers whose corresponding port numbers are not displayed in the search.

- Identify undiscovered servers based on server serial number (Service Tag)

Type `Key = ServiceTag`

Displays service tag for all discovered servers. Servers not displayed here were not discovered.

- Identify undiscovered servers based on server MAC address

Type `Mac address of vmknic =`

Displays the MAC address of all discovered servers. Servers whose MAC addresses are not displayed were not discovered.

Server Imaging Failed at Kickstart Delivery Task

Problem

Server imaging failed at task `Kickstart delivery and wait for Host up`.

Cause

The kickstart delivery task fails if the server failed to request the kickstart script after pulling the ESXi image from VIA. This is usually due to network issues or due to an ESXi installer crash.

Solution

- ◆ Check the server console to see if the server is actually booted into ESXi and shows the expected IP address.

Server Imaging Failed at Check Host Status Task

Problem

Server imaging failed at task `Check host status`.

Cause

This issue occurs because VIA is unable to log in to the server right after ESXi has been installed on it.

Solution

- 1 Click the alarm icon on the top right corner of the window and look for the following alarm:
`Host 192.168.100.XX is different than the one with Mac Addr aa:bb:cc:dd:ee:ff discovered through DHCP`
This alarm indicates that all servers were powered off before imaging
Identify the servers that were not powered off, power them down, and retry imaging on those servers.
- 2 Search for the following strings in the log file.
`Esxi imaging failed on host OR Maximum retries exceeded`
Log in to the server console for additional troubleshooting.

Cloud Foundation Glossary

13

Term	Description
add rack	Configure an additional rack for a Cloud Foundation system.
additional rack	Additional racks (added after the first rack) to a Cloud Foundation system.
bring-up	Initial configuration of a newly deployed Cloud Foundation system. .
Cloud Foundation system.	Set of physical racks managed as a unit by a single SDDC Manager.
first rack	First (primary) rack in the Cloud Foundation system. The management domain is deployed on this rack.
Hardware Management System (HMS)	Manages hosts and switches in the Cloud Foundation system.
host	An imaged server.
imaging	During imaging, SDDC software is pre-configured on a physical rack.
integrated system	System that combines hardware and software. Can be purchased from select VMware partners. The partner images the rack before sending it to the customer site.
inter-rack switches	Connects individual ToR switches with each other to provide connectivity across racks. These switches are required only when you have more than one rack in your Cloud Foundation system, and are placed on the second rack.
Lifecycle Manager (LCM)	Automates patching and upgrading of the software stack.
management domain	Cluster of physical hosts (first four hosts in the physical rack) that house the management component VMs
management host	Standalone ESXi server to host the Windows jump VM used for imaging.
primary host	Host on which VIA deploys the SDDC Manager VMs during imaging, and which bootstraps the first rack during bring-up.
SDDC Manager	Software component that provisions, manages, and monitors the logical and physical resources of a Cloud Foundation system.
SDDC ManagerController VM	Contains the SDDC Managerservices and a shell from which command line tools can be run. This VM exposes the SDDC Manager UI.
SDDC Manager Utility VM	Contains the LCM depot, backup repository containing NSX Manager and host backups, and 2nd DNS instance.
server	Bare metal server in a physical rack. After imaging, it is referred to as a host.
Top of Rack (ToR) switch	Connects servers within a rack through 10Gbps links to the NICs on each server. A Cloud Foundation rack contains two ToR switches connected to each other.

Term	Description
unassigned host	Host in the capacity pool that does not belong to a workload domain.
workload domain	A policy based resource container with specific availability and performance attributes and combining vSphere, vSAN and NSX into single a consumable entity. A workload domain can be created, expanded, and deleted as part of the SDDC lifecycle operations.