VIA User's Guide

VMware EVO SDDC 1.2

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see http://www.vmware.com/support/pubs.

EN-001780-02

vmware®

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

http://www.vmware.com/support/

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

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

Contents

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 *EVO SDDC Overview and Bring-Up Guide* contains detailed information about the EVO SDDC product, its components, and the network topology of an EVO SDDC installation.

The *Administering VMware EVO SDDC* provides information about how to manage a VMware EVO SDDC™ 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.

About VIA

1

VIA is a virtual appliance that enables VMware partners to image physical racks with VMware SDDC software before shipping them to customers.

VMware provides the VIA OVA template and a software bundle to partners. The software bundle consists of key SDDC components such as VMware vSphere ESXi, vSphere, NSX, Virtual SAN and corresponding management tools such as vRealize Operations Manager and vRealize Log Insight. The versions, editions, and patch levels of all products in the software bundle are pre-specified and pre-qualified.

The partner's infrastructure includes an ESXi host (referred to as the management host) and a 24-port 1GE Managed Switch with RJ45 ports and Cat 5/5E cables. VIA uses 3 ports on the managed switch (which allows partners to create VLANs). The partner installs the VIA OVA template on the management host and uploads the software bundle into the VIA VM. The management host is connected to the public network as well as to the private network used by the VIA VM to image the individual hosts and switches. A jump VM provides an interface on the public network for partners to connect to the VIA VM to image the physical rack.

Figure 1-1. VIA Deployment

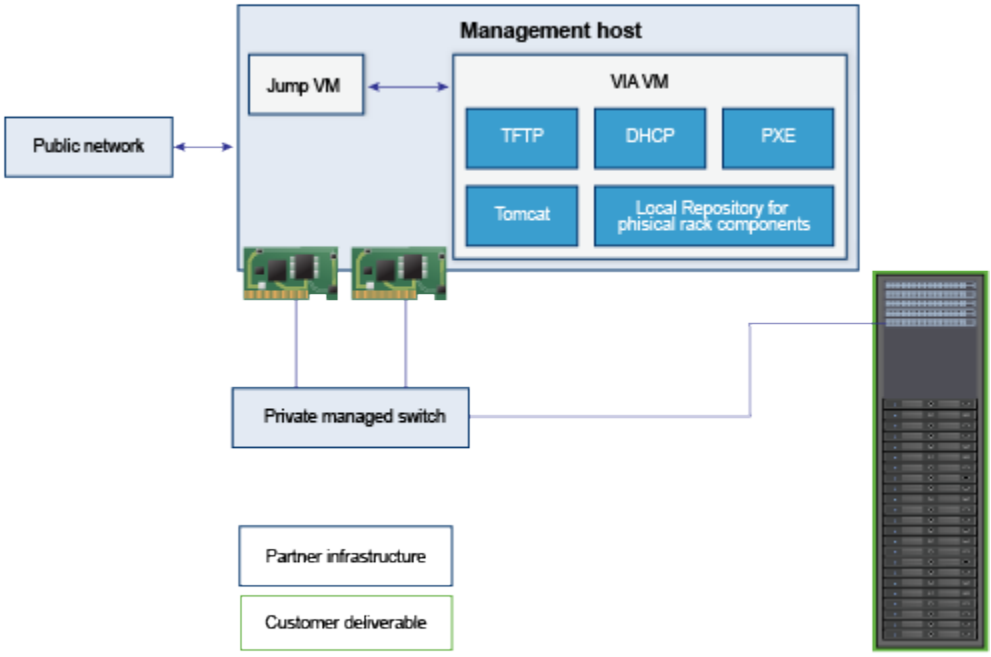

The physical rack consists of a management switch, two spine switches, two Top of Rack (ToR) switches and up to 24 physical servers. The management switch is the first device in the physical rack to be imaged by VIA and provides access to the other devices. The spine switches are imaged next followed by the ToR switches. The servers in the physical rack are then imaged in parallel. The SDDC software is loaded on to the first server (node 0) in the physical rack. 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 shipped to the customer.

This chapter includes the following topics:

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.

The bundle contains the following software:

- ESXi
- vCenter Server
- NSX
- Virtual SAN
- vRealize Log Insight
- vRealize Operations
- EVO SDDC Manager
- Platform Services Controller

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 HSQLDB 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.

Inventory

VIA maintains a bundle inventory and a rack inventory.

The bundle Inventory is an input to the imaging activity, and is created by VIA before it begins an imaging run. The bundle inventory is specific to a vendor and hardware type.

The rack inventory is an enumeration of the configuration details of the hardware imaged by VIA. The configuration details includes credentials to access both the data and the management interfaces of the imaged hardware, as well as the protocols to be used to access the interfaces of the imaged hardware.

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.

Bundle Inventory Service

VIA deploys the bundle inventory service before starting an imaging activity. The service creates a bundle inventory using all the information in the bundle manifest. It ensures that the software bundle contains the files listed in the manifest and lists the manufacturer and hardware for which the bundle can be used.

The bundle inventory service includes a bundle manager and bundle controller. They manage the software bundles, mount the active bundle to be used for an imaging run, and set up TFTP and PXE Linux configuration to image the servers.

Device Manifest Service

The device manifest service creates a new manifest file when an imaging activity is performed for the first time. It also tracks changes to the device status and stores hardware information for the rack components.

Imaging Service

The imaging service can start, stop, or cancel an imaging run. It tracks the imaging workflow and maintains the state of the imaging run as well as details about the device being imaged. Details being tracked include the IP address of the device, imaging task being performed, status of the imaging task, and completion time of the imaging task.

DHCP Service

VIA deploys the DHCP service before starting an imaging run. The DHCP service discovers the physical rack components and their PXE images using the DHCP Protocol. It keeps track of the IP addresses allocated to the devices to ensure that a device can be provided with the same IP address in case it needs to be reimaged.

Cipher Service

The imaging service uses a cipher service to generate passwords to access the imaged rack components. The cipher service ensures that each imaged component is always associated with a unique password. However, all ESXi hosts have the same password.

Rack Inventory Service

The rack inventory service is deployed when the components have been successfully imaged. It collects access information for the imaged components such as connection protocol, IP address, and username and password and generates an inventory file. This inventory file is then transferred to the management switch.

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
Spine switches	NA	2 (Rack 2 only)
TOR Switches	2	2
Management switch	1	1
Servers	Up to 24	Up to 24

EVO SDDC does not come with a console serial switch, but it is a nice addition to your environment.

- PDUs

  Each physical rack must have 4 PDUs (2 primary and 2 standby) even if it contains less than 24 servers. The primary PDUs must be blue and the standby must 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.

- Spine switches

  Rack 2 in your EVO SDDC system contains two 32 x 40 GE spine switches. These switches connect multiple racks by using uplinks from the Top of Rack switches.

  Spine switches should not be connected to a public 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 the primary rack connect EVO SDDC to the public network.

- Management switch

Each rack contains a 1 GE management switch, which is used for IPMI access and access to the physical switches. The management ports of the ToR switches, Spine switches, and the physical servers are connected to the management switch. The data ports of the ToR switches are also connected to the management switch. This enables the management switch to monitor the data from the servers from both the management network as well as the data network.

The management switch provides out-of-band (OOB) connectivity for managing switches and servers. The hardware management service (HMS) runs on the management switch.

- Servers

  A rack can contain up to 24 two-socket 1U servers including the latest Intel Xeon processors with high-performance DDR4 memory and 10 storage device slots. Each server includes an embedded NIC with 2 x 10 GE interfaces. Each 10 GE interface is connected to a separate ToR switch to ensure redundancy. A server's out-of-band management interface is connected to the management switch by using a 1 GE interface.

  Each server contains the following:

  - 2 Intel E5-2600 series v3 CPUs
  - 384 GE DDR4 ECC
  - 2 x 10 GE ethernet network interfaces
  - 1 GE out-of-band management interface
  - Upto 10 storage slots with 8 x HDDs and 2 x SSDs

Figure 1-2. Example Physical Rack Configuration

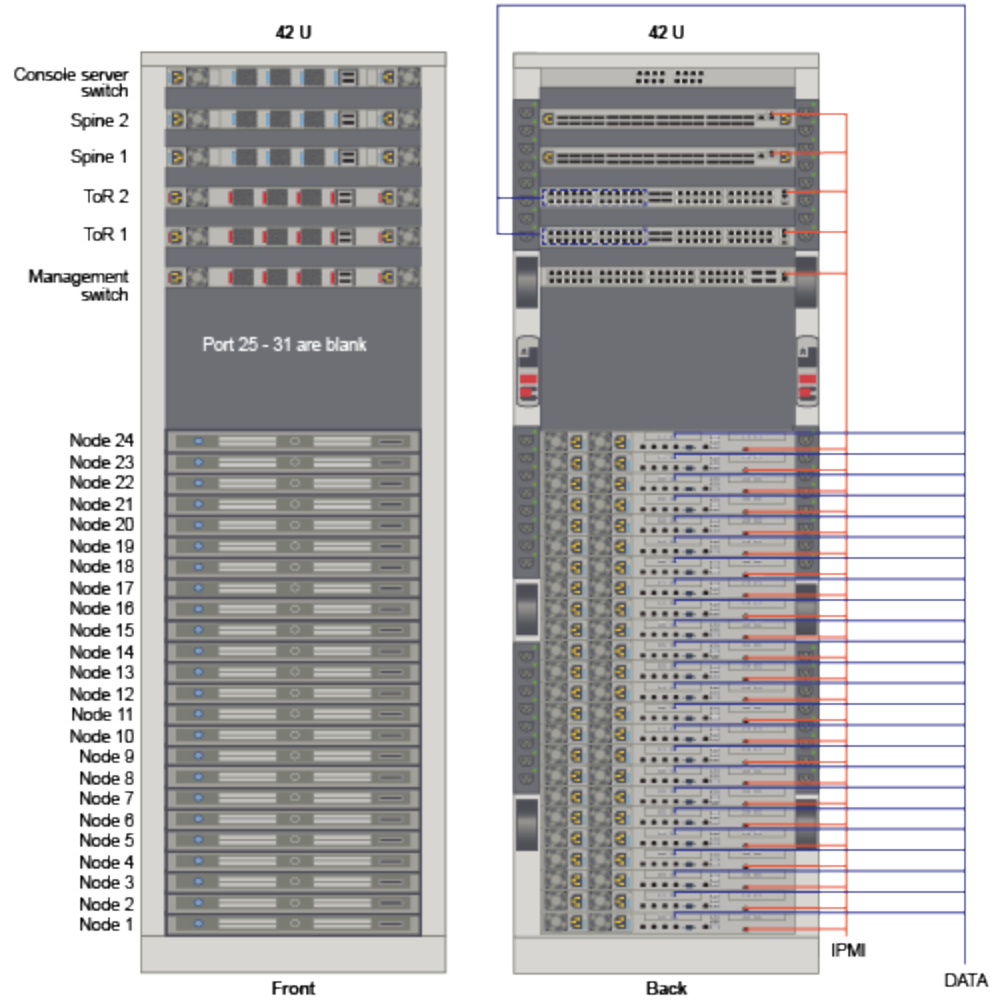

NOTE The above graphic is being drawn to scale.

Before You Install VIA 2

Before you install VIA, ensure that you have all of the required hardware components in place.

This chapter includes the following topics:

Requirements for VIA

VIA requires the following infrastructure.

- Management host - a standalone VMware vSphere ESXi 5.5 or later server to host the Windows jump VM. The management host must have at least two NICs, with one NIC connected to the public network and one NIC connected to the private network.
- Jump VM to access VIA
- Private managed switch. Private indicates that only you are using it. A 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

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

Review the Bill Of Materials (BOM) from VMware and ensure there are no discrepancies between the BOM and the equipment being used.

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 Spine 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.

Regardless of the number of servers in each rack, cables must be in place for 24 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
- Managment switch ports 49 and 50 going to the ToRs: black
- ToR 1 cables to servers: blue
- ToR2 cables to servers: red
- ToR 1 and ToR 2 connections to spine 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.

Ensure that the cable type and length being used in your setup meet the following requirements.

- The cable connecting the physical server baseboard management controller (BMC) port to the management switch is 10 ft.
- The cable connecting the physical server 10 G interfaces to the ToR switches is 1-2 m (3.28-6.56 ft).
- The cable connecting the ToR switches 40G interfaces to the Spine switches is 1-2 m (3.28-6.56 ft).

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 velco 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.

- Test cables from the physical server baseboard management controller (BMC) port and the management switch.
- Review the test print out to confirm that the cables passed the test.
- Cables from the physical server BMC port to the management switch must be seated properly.
- 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 Spine switch.

- There must be two 40G QSFP+ cable connections between eachToR switch and Spine switch.
- Inter-switch SFP+ and QSFP+ cables must be seated properly.

Rack Wiring

Connect cables according to the wiremaps.

Rack 1

Figure 2-1. Wiremap for rack 1 with Dell Components

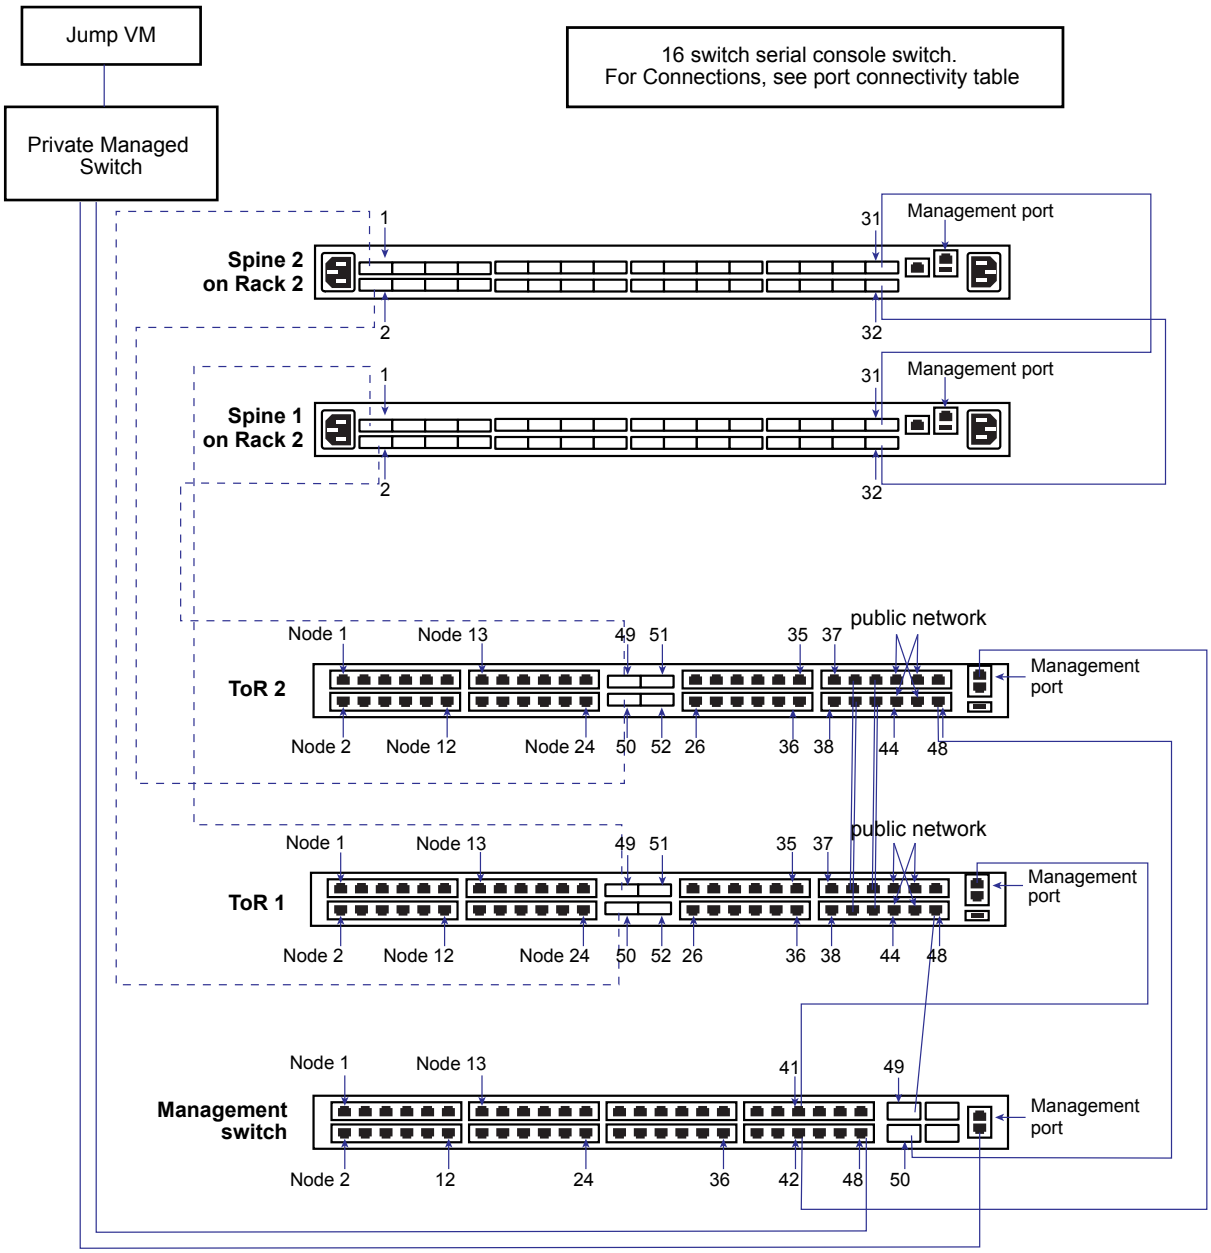

Figure 2-2. Wiremap for rack 2 with Dell Components

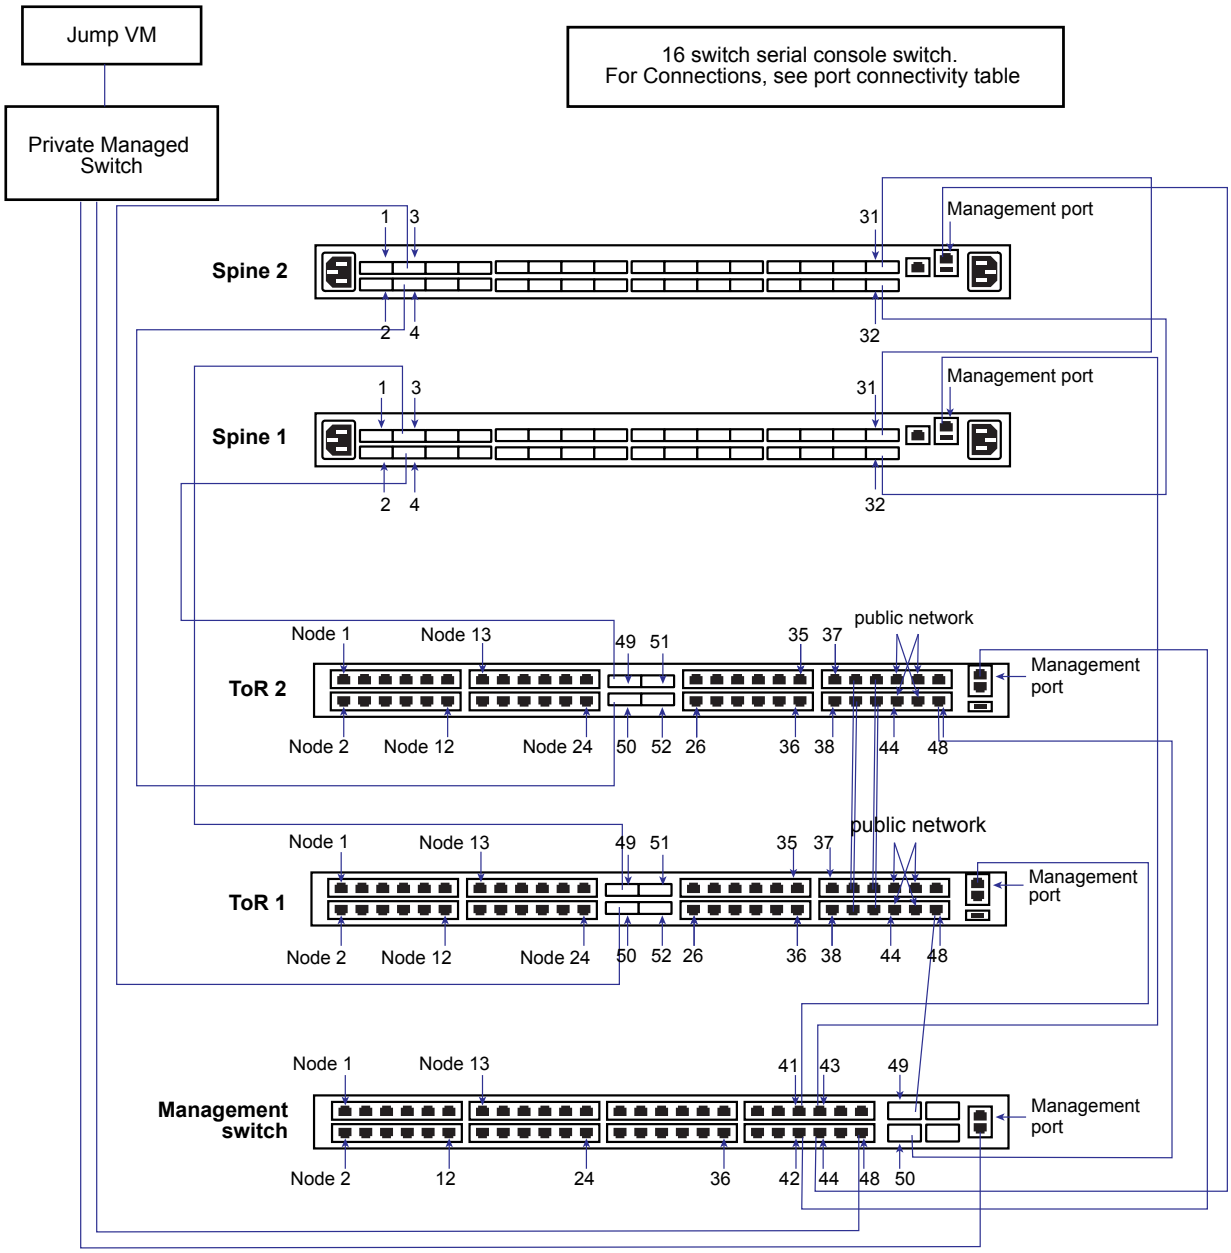

Figure 2-3. Wiremap for rack 1 with Quanta Components

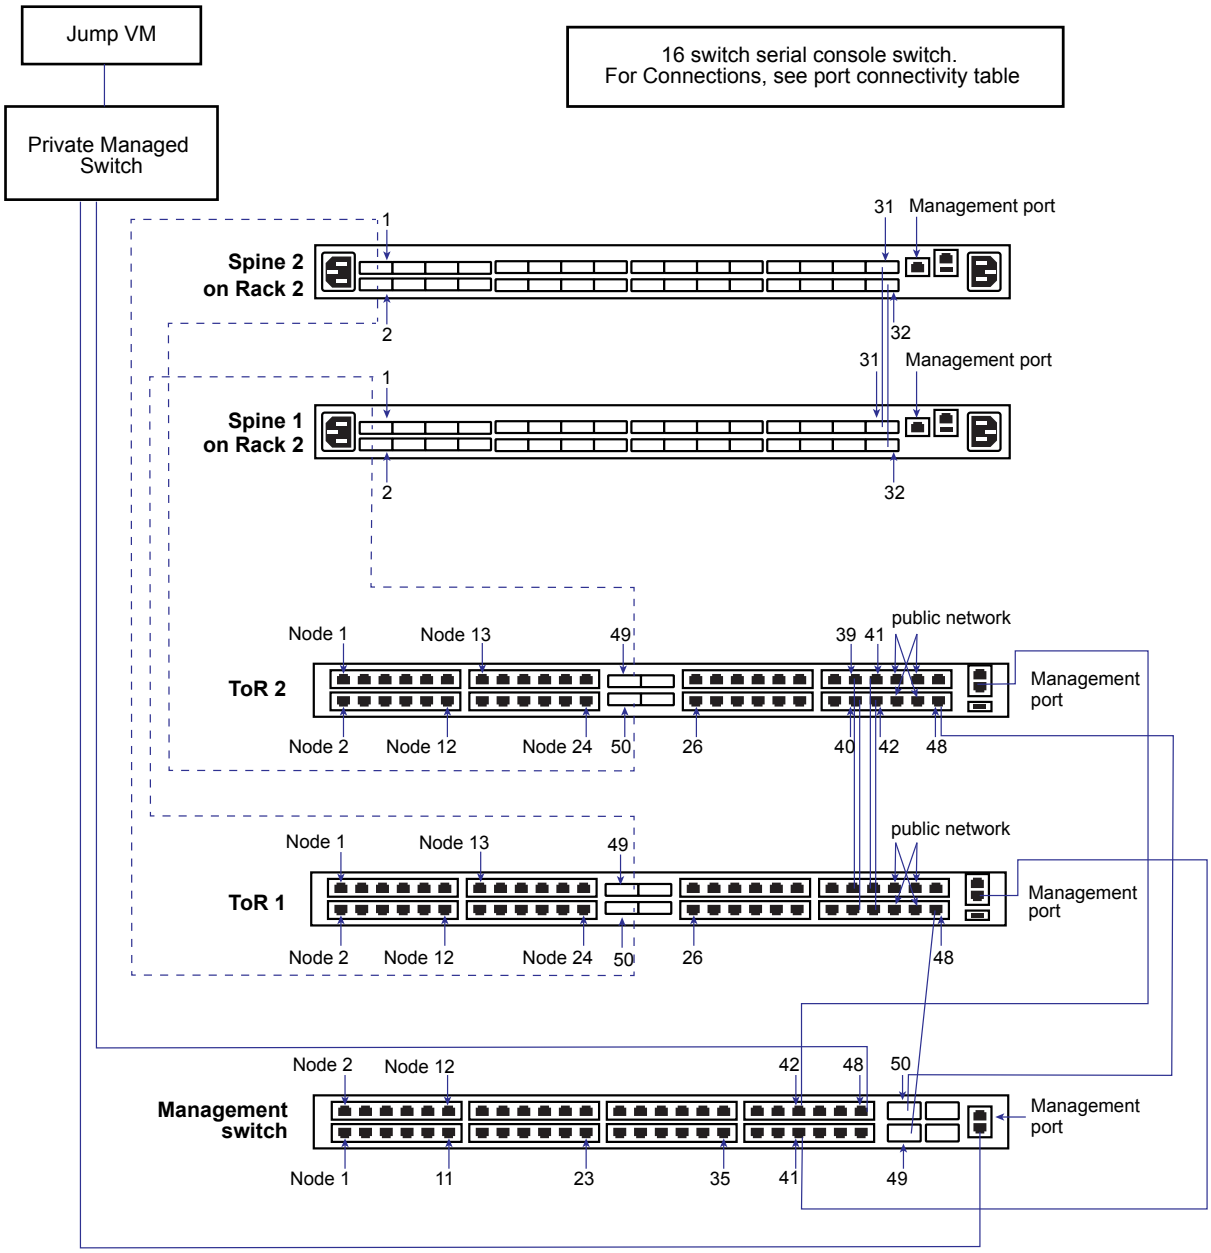

Figure 2-4. Wiremap for rack 2 with Quanta Components

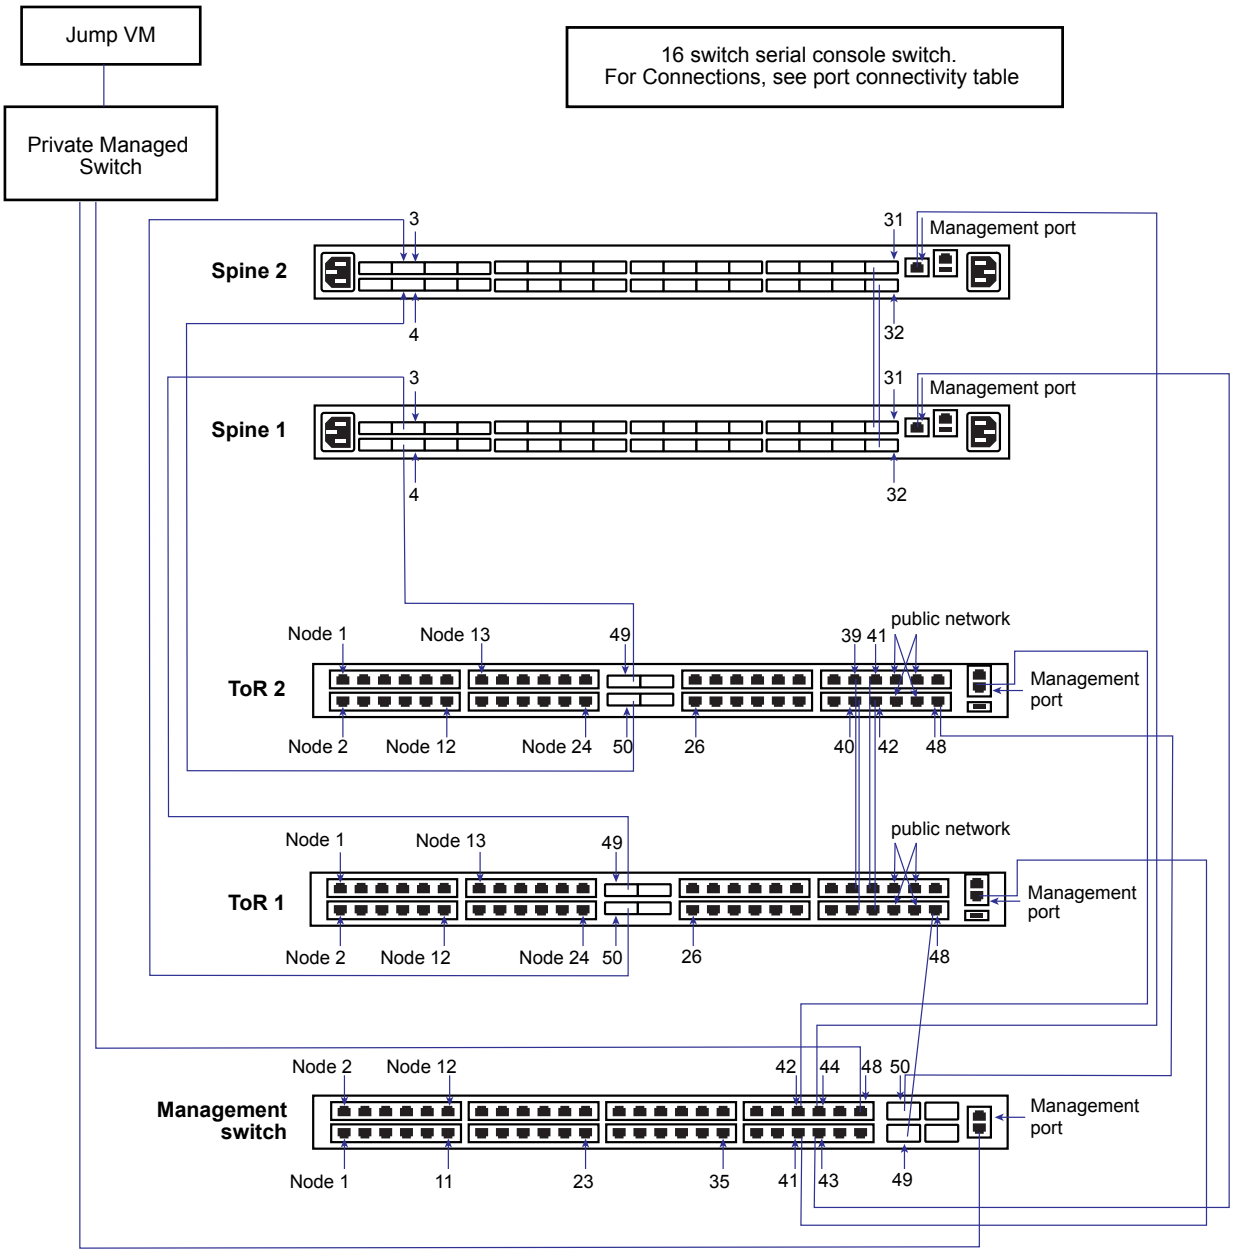

Additional Racks

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

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

Rack Component Ports

Refer to the tables below for port connectivity information.

Console Serial Switch

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

Spine 2 (Rack 2 only)

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

Spine 1 (Rack 2 only)

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

Port Number	Speed	Connects To
3	40 Gbps	Rack 1 ToR 1 port 49
4	40 Gbps	Rack 1 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 1 port 49
11	40 Gbps	Rack 6 ToR 1 port 49
12	40 Gbps	Rack 6 ToR 1 port 49
13	40 Gbps	Rack 7 ToR 1 port 49
14	40 Gbps	Rack 7 ToR 1 port 49
15	40 Gbps	Rack 8 ToR 1 port 49
16	40 Gbps	Rack 8 ToR 1 port 49

ToR 2

Port Number	Speed	Connects To
1 - 24	10 Gbps	node 1 - node 24 where port 1 connects to node 1, port 2 connects to node 2, and so on
25-38	NA	Not connected
39-42	10 Gbps	ToR 1 ports 39 - 42
43-46	10 Gbps	Public network
47	NA	Not connected
48	1Gbps	Management switch port 50
49	40 Gbps	Spine 1 port 2
50	40 Gbps	Spine 2 port 2
51-52	NA	Not connected
Management	1 Gbps	Management switch port 42

ToR 1

Port Number	Speed	Connects To
1 - 24	10 Gbps	Node 1 - node 24 where port 1 connects to node 1, port 2 connects to node 2, and so on
25-38	NA	Not connected
39-42	10 Gbps	ToR 2 ports 39 - 42
43-46	10 Gbps	Public network
47	NA	Not connected
48	1Gbps	Management switch port 49
49	40 Gbps	Spine 1 port 1

Port Number	Speed	Connects To
50	40 Gbps	Spine 2 port 1
51-52	NA	Not connected
Management	1 Gbps	Management switch port 41

Management Switch

Port Number	Speed	Connects To
1 - 24	1 Gbps	Node 1 - Node 24 where port 1 connects to node 1, port 2 connects to node 2, and so on
25 - 40	NA	Not connected
41	1Gbps	ToR 1 management port
42	1Gbps	ToR 2 management port
43	1Gbps	Spine 1 management port
44	1Gbps	Spine 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 Rack Unit (RU) location of each device.

Hardware Devices

Table 2-1. Physical Device Location in Primary Rack

RU Location	Device
42	Spine 2
41	Blank
40	Spine 1
39	Blank
38	ToR 2
37	Blank
36	ToR 1
35	Blank
34	Management switch
25-33	Blank
1-24	Nodes 1-24

Table 2-2. Physical Device Location in Additional Racks

RU Location	Device
39-42	Blank
38	ToR 2
37	Blank
36	ToR 1
35	Blank
34	Management switch
25 - 33	Blank
1-24	Nodes 1-24

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.

BIOS Settings

The Bill of Materials (BOM) specifies the BIOS settings for each device. Ensure that the settings on the physical devices in your environment match the BIOS settings in the BOM.

Firmware Settings

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

Network Switches

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.

ONIE version

Ensure that the correct ONIE version is installed as per the BOM.

Management Host

Physical Connectivity

The management host and 24-port private managed switch are located at the partner site.

Table 2-3. 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 2-5. Management Host Connection

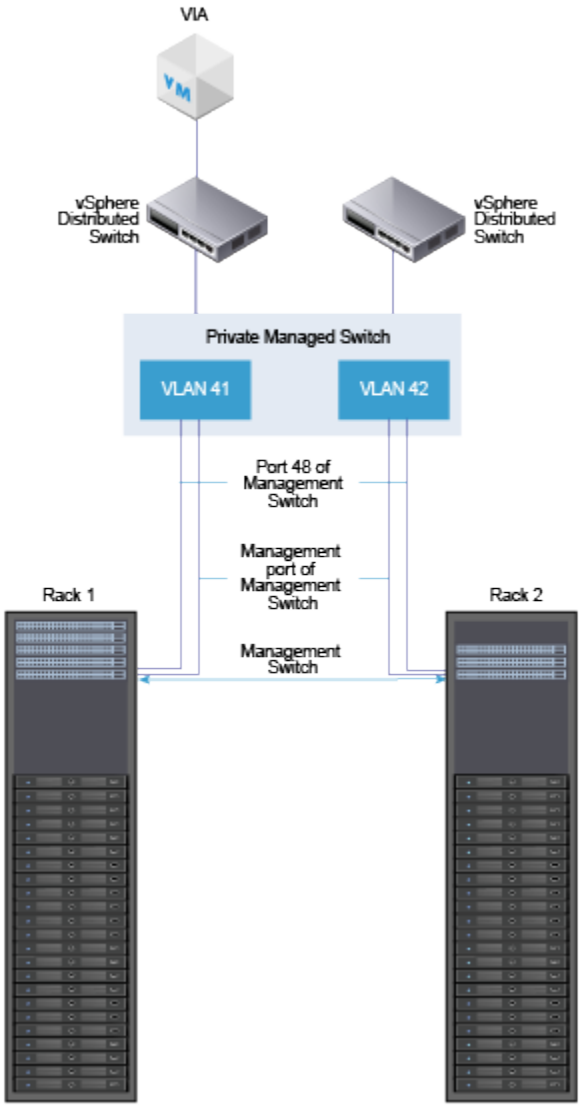

Private Managed Switch

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 version 5.5 or later on the local disk.
- 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 `0.vmware.pool.ntp.org`.

  It is important to ensure that the time on the management host is set correctly.
- Enable SSH on the managament 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 following virtual machines run on the management host.

- A VIA VM
- A jump VM

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 2-4. Jump VM Hardware Configuration

Virtual Hardware	Value
Memory	4 GB
vCPU	1 virtual socket, 2 cores per socket
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
Navigate to **Options > Advanced > General**	Disable logging keyboard.typematicMinDelay = "2000000"

Software Configuration

Perform the following tasks to prepare the jump VM.

- Install the Windows 2012 Essentials operating system on the VM .
- 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 5.5 or later Client
  - VMware Ruby vSphere Console (RVC)
  - 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.
- Add a route to allow BMC access to the physical servers. For example,

  `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, follow the steps below.

  a In a command window, type the command **netsh**.

  b Type the command `int ipv4 show interfaces`.

Pre-Imaging Checklist

Partners 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.

Table 2-5. Pre-Imaging Checklist

Setting	Verified
Review the Bill of Materials (BOM) from VMware and verify that there are no discrepancies between the BOM and the hardware being used. If there is a discrepancy, contact VMware Support.	
Validate that BIOS Settings for all components are correct. See Chapter 7, "BIOS Settings," on page 51.	
Ensure that the correct ONIE version is installed as per the BoM.	
Verify that firmware settings are set correctly as per the BoM.	
Connect each device in the rack to both PDUs.	
Keep power and network cable lengths to a minimum.	
Use specific colors for cables from each device. See "Network Cables," on page 14.	
Verify that the cable bend radius is proportionate to the external diameter. See "Network Cables," on page 14.	
Verify that cables are properly routed and labelled.	
Test cables to ensure that installed cabling links provide the transmission capability to support the required data communication.	
Verify that the physical racks are wired according to the wiremap. See "Rack Wiring," on page 16.	
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.	
Verify that switches have redundant power supplies and each power supply is connected to a separate power strip.	
Ensure that ESXi version 5.5 or later is installed.	
Verify that the **Allow virtual machines to start and stop automatically with the system** option is enabled.	
Assign IP address 10.1.0.200 to the vmk0 kernel interface.	
Verify that SSH is enabled on the management host.	
Verify that the access VM, VIA VM, and jump VM meet the required hardware configuration. See "Virtual Machines," on page 25.	
Verify that the required software has been installed on the VMs. See "Virtual Machines," on page 25.	

Installing VIA

3

VIA is a virtual machine appliance. You need a DCPN account to download the VIA OVF template and software bundle from sftp2.vmware.com. After you install the VIA VM and configure a jump VM on the management host, you copy the software bundle to the VIA VM. You can then access the VIA user interface through a browser on the jump VM.

Prerequisites

- Ensure that you have the infrastructure for VIA available and that you have set up your physical environment as described in Chapter 2, "Before You Install VIA," on page 13.
- Ensure that you have copied the EVO SDDC software bundle to the management host.
- Download the VIA OVF file and bundle ISO image on your local file system.

Procedure

1 Deploy the VIA OVF file using vSphere Client in an isolated (private network). VIA must not be able to access the public network.

   a Login to the vSphere 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 EVO SDDC 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 EVO SDDC bundle file to this folder.

3 Configure time settings on the management host.

   a In the vSphere 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 public 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. Verify the address that you want to use against the `via.properties` file in the bundle ISO to avoid any conflict.

6 Download 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 Power on the VIA VM.

   VIA is deployed with pre-configured network settings and is available at IP address 192.168.100.2. 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.

9 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.

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

4

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

In a multi-rack environment, you can either image all racks in parallel, or image the primary rack 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 4-1. VIA Setup for Parallel Imaging of Multiple Physical Racks

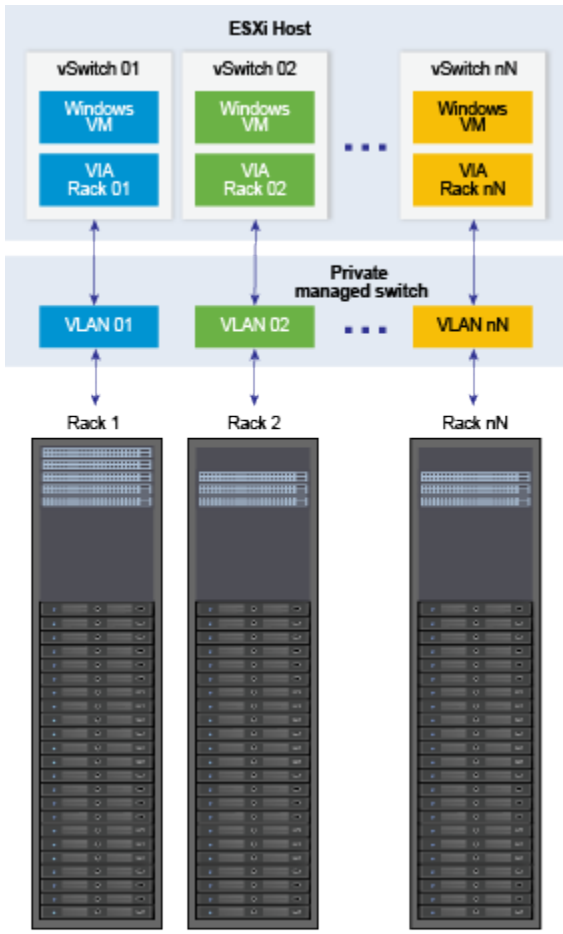

This chapter includes the following topics:

- "Resume Imaging," on page 42
- "Image Additional Racks," on page 44

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 EVO SDDC management data is routed. The management ports of the ToR switches, Spine switches, and the physical servers are connected to the management switch. The data ports of the ToR switches are also connected to the management switch. This enables VIA to communicate with the servers over both management and data network through the management switch. VIA is also 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. VIA currently uses an IPMI connection to image the management switch.

2 Spine switches and ToR switches

Spine and ToR switches are imaged in parallel.

Spine switches inter-connect multiple racks enabling a scale out architecture for the datacenter. They create an stretched L2 backplane between racks.

ToR switches provide connectivity to servers in each rack out to spine switches. The first pair of ToR switches provide connectivity to your datacenter network.

3 Servers

The management ports on the servers become accessible to the management switch during the course of imaging/configuration, which in turn make the management ports accessible to VIA through the management switch. Once all the switches are imaged and configured, the data ports of the servers become accessible to VIA through the ToR switches, which then proceeds to image the servers in parallel.

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

1 Discovery

Rack components are discovered using the DHCP service. 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 Spine 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 Spine switches are discovered and imaged. The management switch also discovers the IPMI network of the servers. This allows VIA to initiate imaging of the servers. 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 spine switches, imaging stops at that point and cannot proceed. If a configuration step fails for the server, imaging for that server cannot be completed but the remaining servers in the rack can be imaged.

Table 4-1. Management Switch Configuration

Number	Step Name	Description
1	Apply license	Apply the relevant license to the installed image
2	Configure ports	1 Configure the ports which allow the management switch to connect to the management interfaces of the ToR and spine switches and the servers. 2 Bridge the ports connected to VIA with the ports connected to the management interfaces of the ToR and spine switches. 3 Create separate subnets for the management network and data network of the rack.
3	Update interface	Ensure that only the management interfaces of ToR and spine switches are enabled while the management interfaces of the servers and the data network interfaces of the ToR switches are disabled before initiating the imaging of ToR switches.
4	Setup persistent network	1 Wait till all ToRs are imaged and bridge the ports connected to VIA with the ports connected to the management interfaces of the ToR and spine switches and servers to enable VIA to listen to DHCP requests. 2 Setup Spanning Tree Protocol (STP) on the IPMI management interfaces of the ToR and spine switches and servers. 3 Setup STP on the EVO SDDC management interfaces of the ToR switches and the interfaces connected to VIA. 4 Enable LACP on ToR data interfaces. 5 Create separate subnets for the management network and data network for the rack.
5	Setup IPMI DHCP	Set up a DHCP service to discover the IPMI network of the servers.
6	Host Power Cycle	Discover all servers and ensure that the minimum required servers are available to ensure that EVO SDDC can be deployed. If the requirement is met, VIA initiates a power cycle of all servers to initiate their imaging. If the required number of servers are not detected, imaging is aborted.
7	Change Password	Change the default password to connect to the switch and stores the new password in a password store.
8	Generate Manifest	Generate device manifest, which contains the current state of the imaging activity for each rack component.

Table 4-2. ToR Switch Configuration

Number	Step	Description
1	Apply license	Apply the relevant license to the installed image.
2	Configure ports	Configure all ports on the switch to operate in Full Duplex mode with auto negotiation enabled and at 1000Mb/s.
6	Change password	Change the default password to connect to the switch and stores the new password in a password store.
7	Generate Manifest	Generate device manifest, which contains the current state of the imaging activity for each rack component.

Table 4-3. Spine Switch Configuration

Number	Step	Description
1	Apply license	Apply the relevant license to the installed image.
5	Change password	Change the default password to connect to the switch and stores the new password in a password store.
6	Generate Manifest	Generate device manifest, which contains the current state of the imaging activity for each rack component.

Table 4-4. Node 0 Configuration

Number	Step
1	Wait for kickstart delivery.
2	Check host status.
3	Install VIBs.
4	Run storage configuration script.
5	Check VSAN setup.
6	Reboot host.
7	Post- ESXi installation configuration.
8	Verify disk status.
9	Create user task.
10	Check VSAN status after reboot.
11	Deploy LCM.
12	Shutdown LCM.
13	Take LCM snapshot.
14	Deploy LCM backup VM.
15.	Shutdown backup LCM VM.
16	Take backup LCM snapshot.
17	Deploy ISVMs.
18	Shutdown ISVMs.
19	Take ISVM snapshot.
20	Deploy VRM.
21	Post VRM installation configuration.
22	Set VM startup shutdown order.
23	Upload bundle ISO.
24	Add ISO to VRM.
25	Collect inventory.
26	Import SSH public keys.
27	Copy PRM manifest.
28	Copy HMS IB inventory.
29	Create VRM snapshot.
30	Reboot VRM.

Table 4-5. Configuration on Remaining Nodes

Number	Step	Description
1	Install Custom VIBs	Install any custom VIBs that may be necessary to enable vendor specific devices on the server.
2	Reboot server	Reboot the server to complete the installation process.
3	Apply licence	Apply ESXi licence.
4	Create user	Create a new ESXi user named EVOSDDC with Administrator role.
5	Generate manifest	Generate device manifest, which contains the imaging status of the device, the IP address assigned, the software used to image it, etc. This is performed on all components irrespective of whether the previous steps were successful or not. This allows VIA to track the status of imaging of any given component during any stage of the imaging process.

Imaging is a multistep process.

1 Upload the Software Bundle on page 35

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 Specify Imaging Details on page 37

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.

3 Monitor Imaging on page 39

In the Monitor Imaging step of the imaging workflow, you can see the imaging status on all devices in your physical rack.

4 Verify Inventory on page 40

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

5 Post Imaging Checks on page 41

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

Upload the 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:

- ESXi
- vCenter Server
- NSX
- Virtual SAN
- vRealize Log Insight
- vRealize Operations
- EVO SDDC Manager
- Platform Services Controller

Prerequisites

- Insert the software bundle CD.

- Download the md5sum file on the jump VM.

Procedure

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

2 Click **Bundle**.

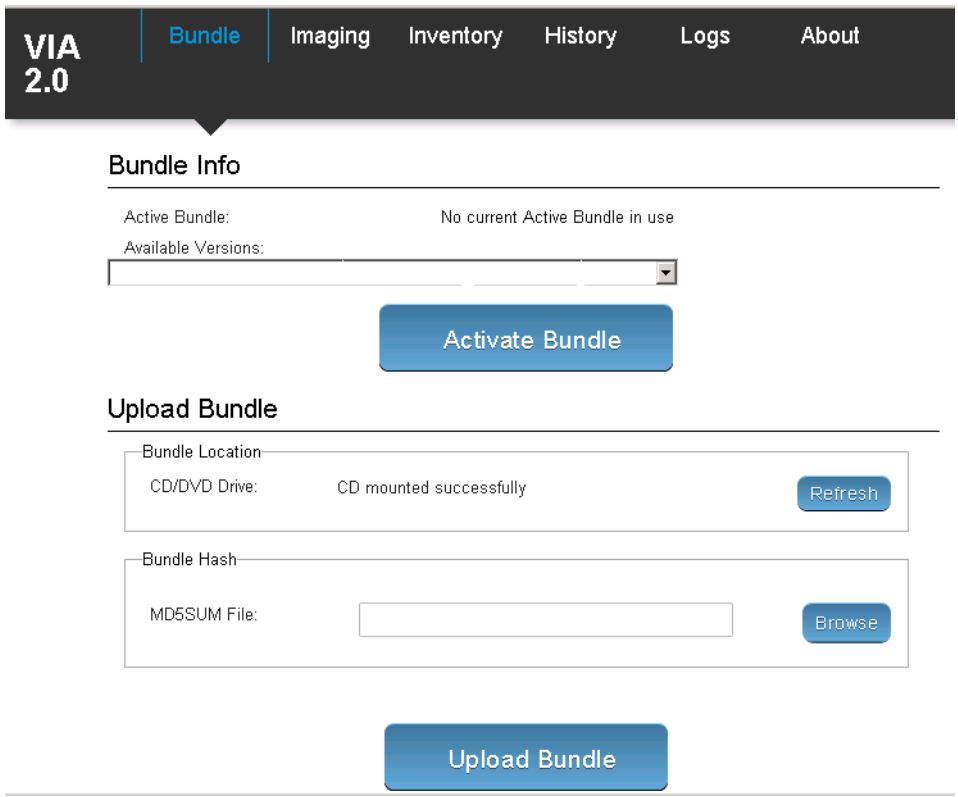

3 In the **Bundle Location** area, click **Refresh**.

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

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

5 Click **Upload Bundle**.

The bundle upload can take several minutes.

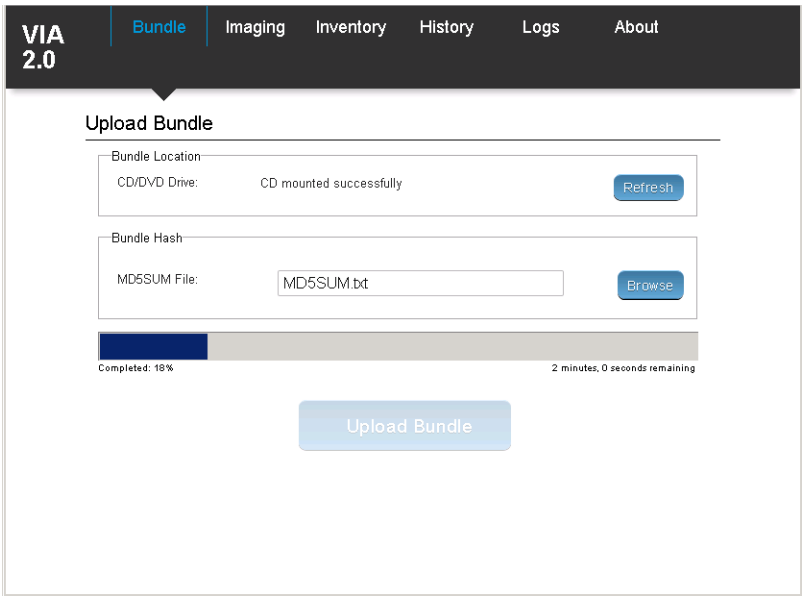

6 In the **Bundle Info** area, select the bundle in **Available Versions** and click **Activate Bundle**.

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**.

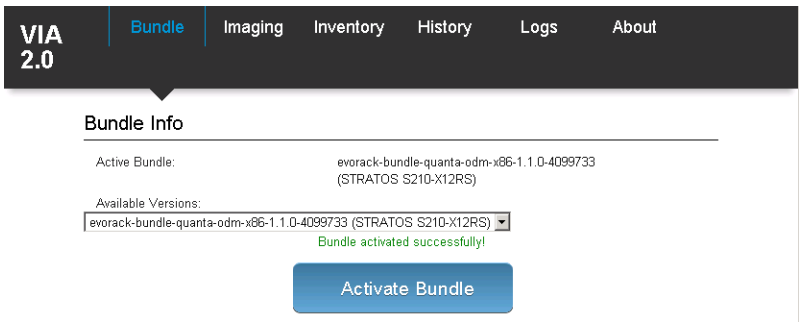

7 (Optional) Verify that the ISO file and `manifest.xml` file are copied to the VIA VM.

a In a console window, SSH to the VIA VM.

```
ssh root@192.168.100.2
```

The password is `root123`.

b Navigate to the `/mnt/cdrom/` directory.

c Confirm that the bundle directory and `manifest.xml` are in this directory.

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.

Procedure

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

Ensure that you are in the Details tab.

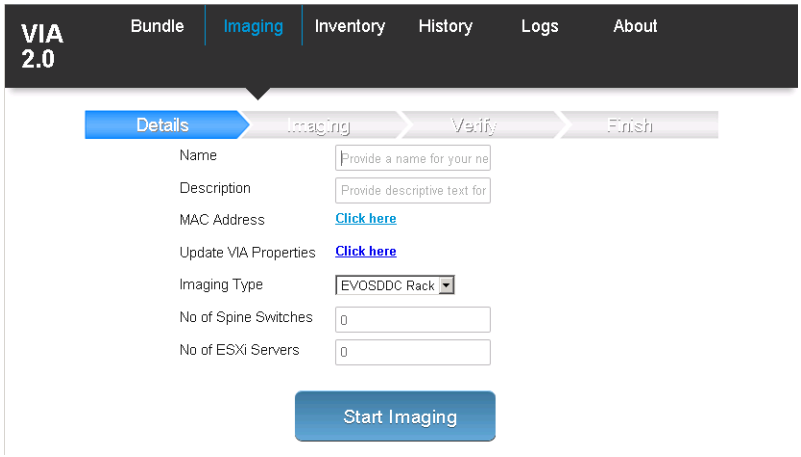

2 (Optional) Type a name and description for the imaging run.

3 Ignore the **MAC Address** field.

4 Click **Update VIA Properties**.

The via.properties file displays rack specification values from the activated software bundle. If required, edit the file as appropriate.

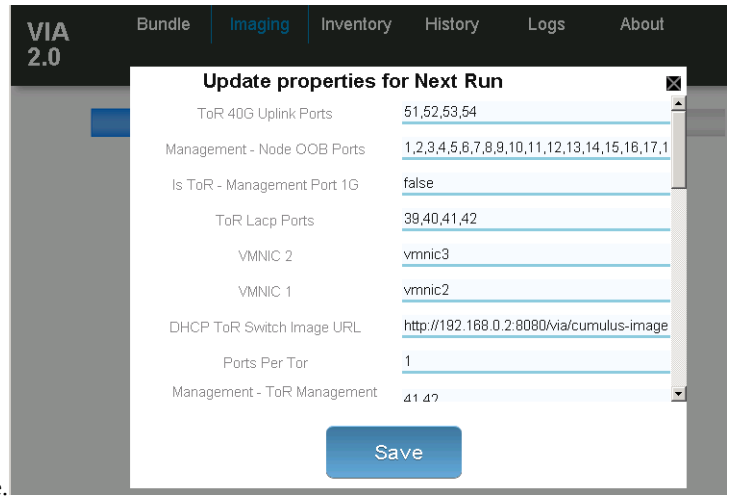

5 Click **Save**.

6 In **Imaging Type**, select **EVOSDDC Rack**.

7 Type the number of spine switches and ESXi servers in the rack you are imaging.

Note Ensure that you type the correct number of spine and ESXi servers to avoid inventory verification failure.

---

8 Click **Start Imaging**.

The **Imaging > Imaging** tab is displayed.

Monitor Imaging

In the Monitor Imaging step of the imaging workflow, you can see the imaging status on all devices in your physical rack.

The **Imaging > Imaging** tab displays the run details, rack details, and imaging status for the rack. The devices in the physical rack are displayed in the order in which they will be imaged.

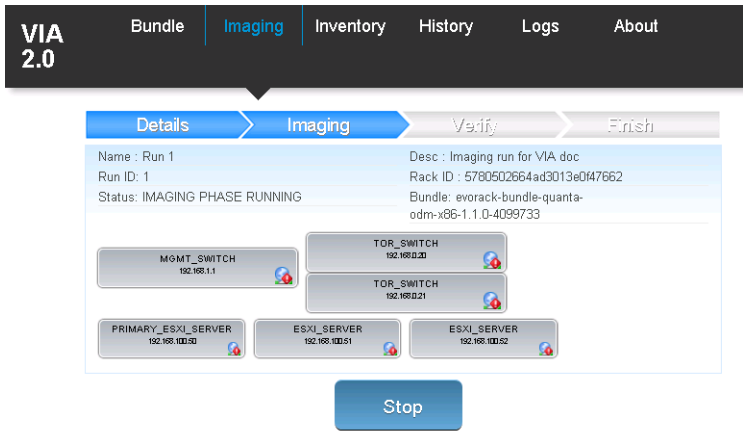

Procedure

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

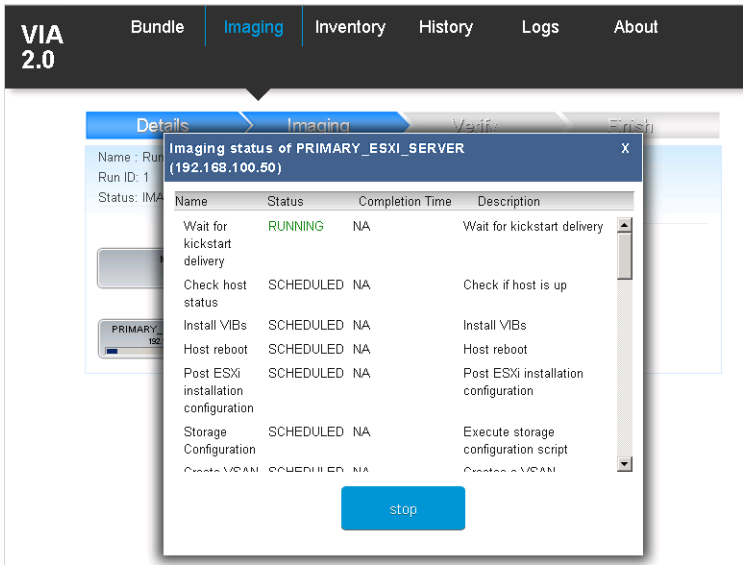

It can take approximately 95 minutes for rack 1 to be imaged. After the imaging is completed successfully, the **Imaging > Verify** tab is displayed.

Note During imaging, the password of all rack components except EVO SDDC Manager is set to `EvoSddc!2016`. The EVO SDDC Manager password is set to a random string, which can be retrieved by an API call.

For information on next steps if a device fails to be imaged, see "Resume Imaging," on page 42.

Verify Inventory

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

The **Imaging > Verify** tab displays the status of inventory collection on each device in the rack.

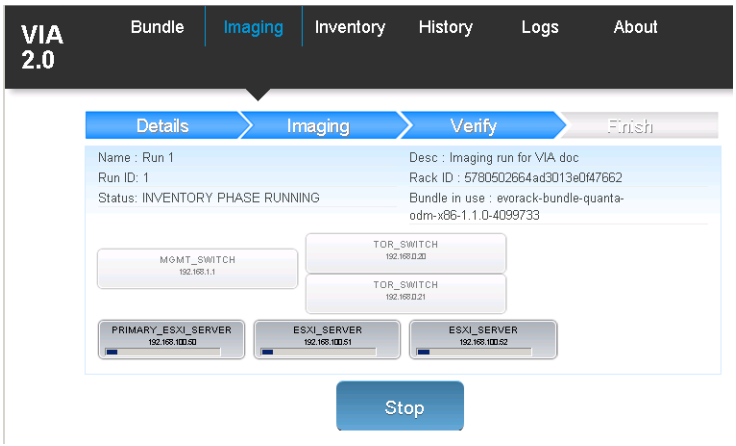

Procedure

- Click a device to see its inventory information. You can expand a component to see more details.

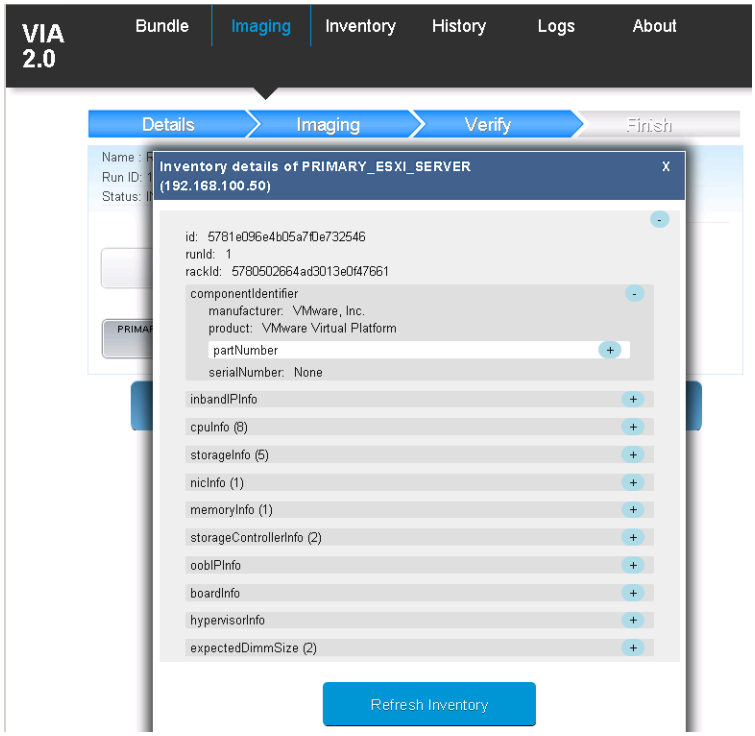

After inventory information for each device has been collected, the **Imaging > Finish** tab is displayed.

Post Imaging Checks

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

After inventory information has been collected for all rack components, the **Imaging > Finish** tab displays

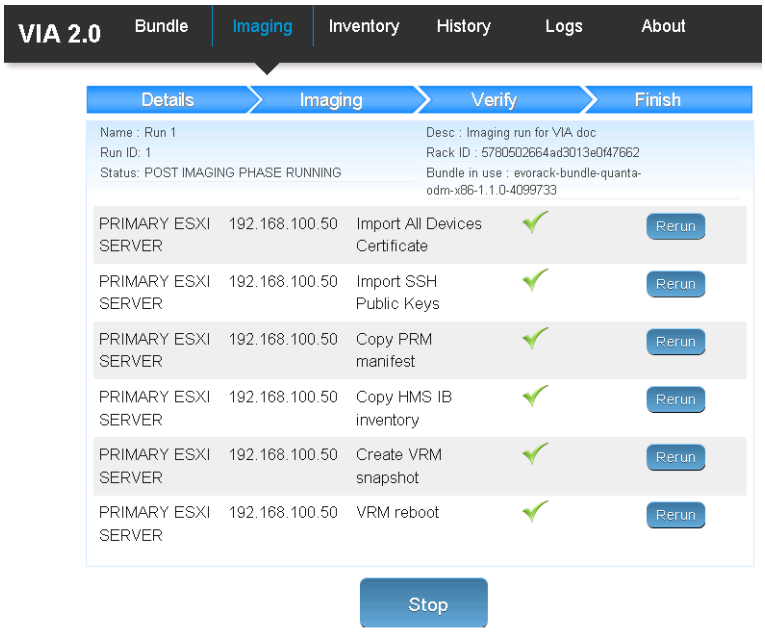

post imaging tasks.

Procedure

1 If a task is not completed successfully, click **Rerun**.

2 After each displayed task has an  icon next to it, click **Finish**.

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

3 Power down the primary rack.

Retrieve EVO SDDC Manager Password and Rack Thumbprint

During imaging, VIA generates a password for the root account of EVO SDDC Manager and a thumbprint for the imaged rack. Both of these are required by the customer.

Procedure

1 Open a new tab in the browser where you were imaging the rack.

2 Type the following:

`192.168.100.2:8080/via/ipsecThumbprint/`*runId*

The browser displays the EVO SDDC Manager password, bootstrap password, and rack thumbprint.

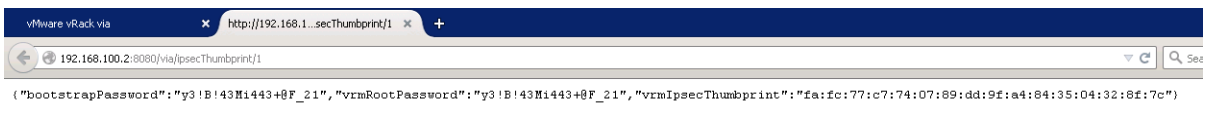

3 Print the output to deliver to the customer along with the imaged rack.

Resume Imaging

If a device fails to be imaged, you can take a number of actions that can help in continuing with the imaging run.

Fix Issues During the Monitor 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.

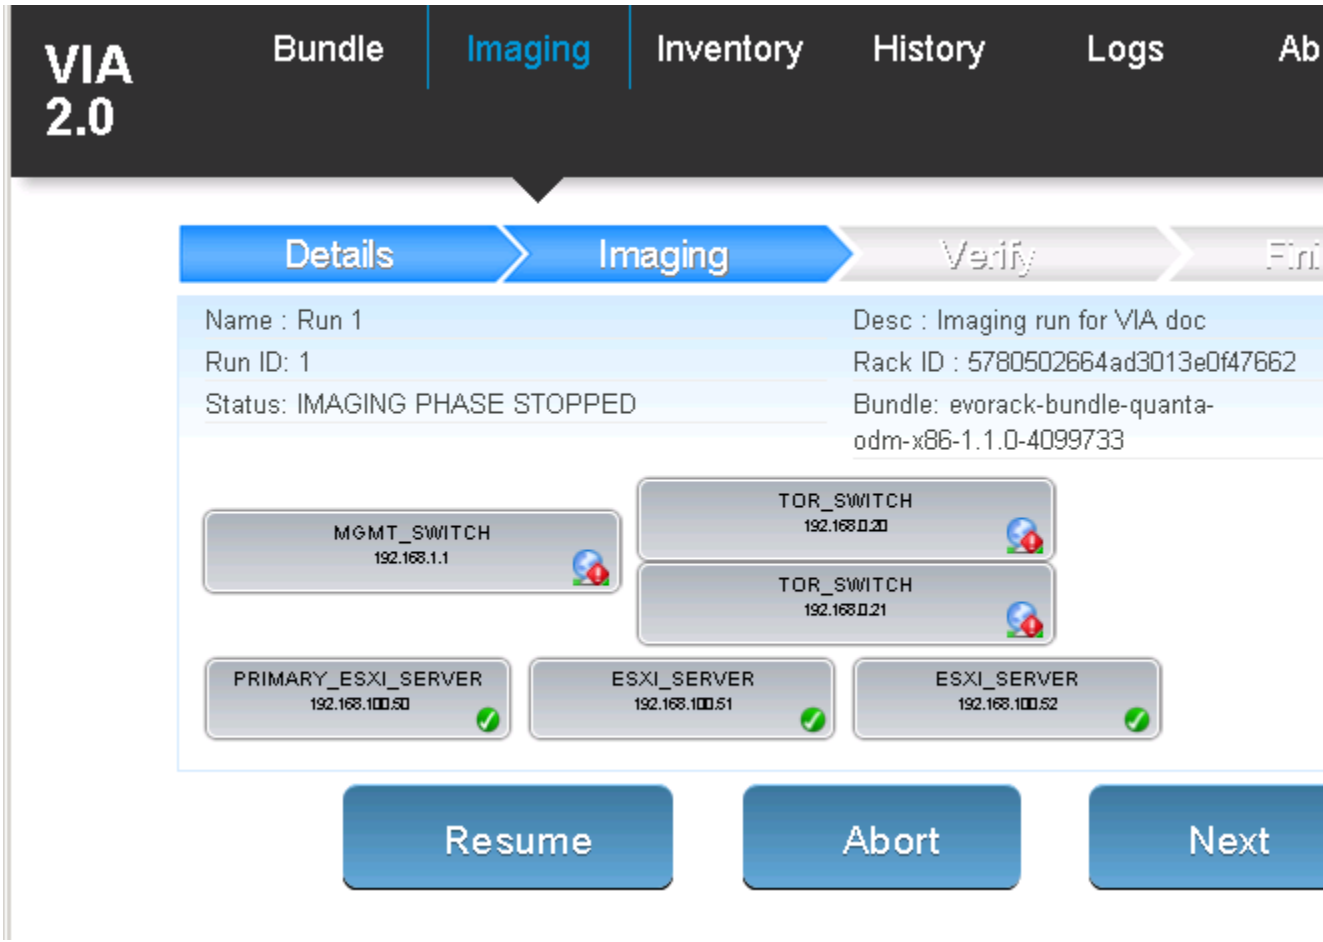

An icon indicates that it has been imaged successfully. An icon indicates that one or more imaging tasks on that devices failed.

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.

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.

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.

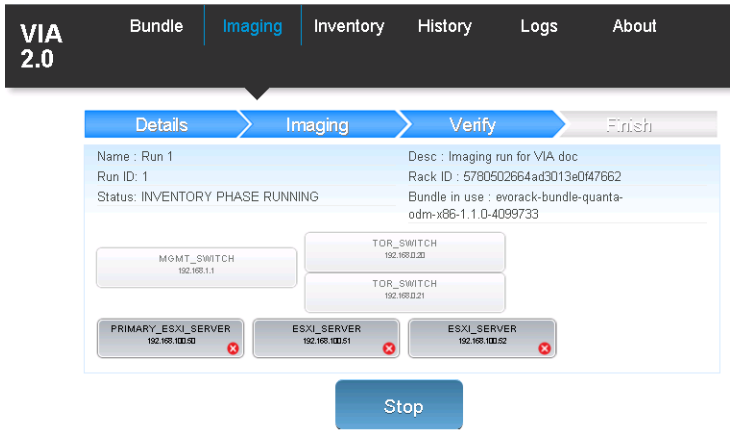

1 Click the component to display the verification task list for that device.

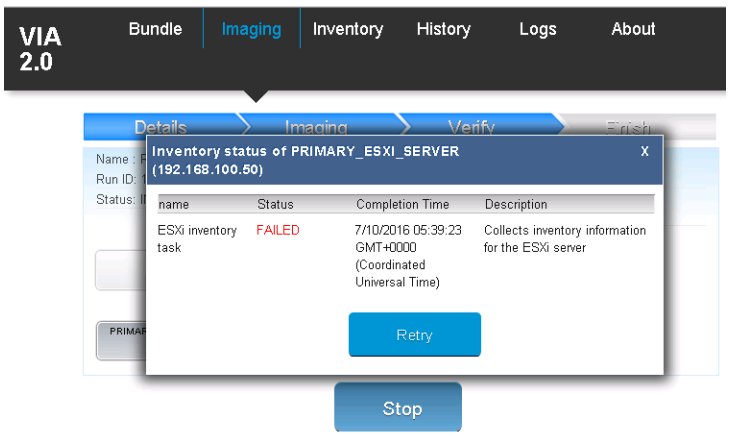

2 Click **Retry**

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 icon.

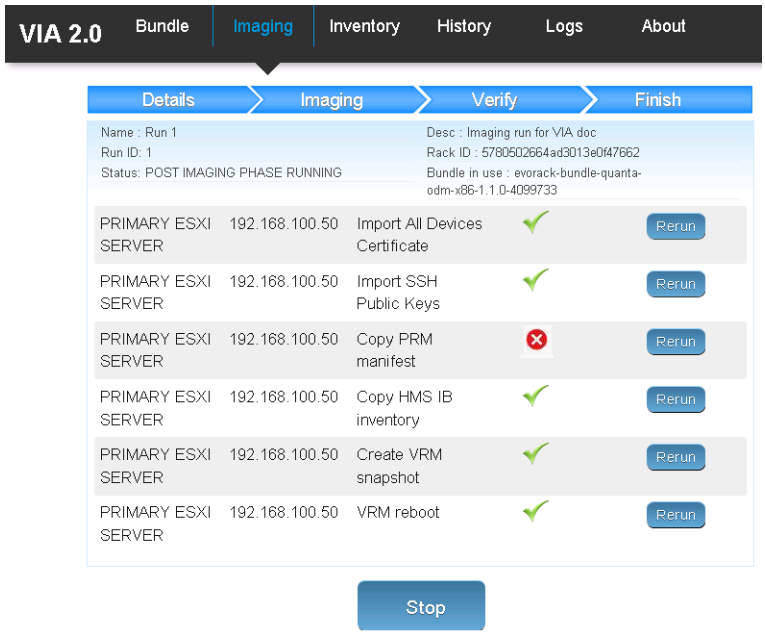

1 Click **Rerun** to run the failed task again.

2 After all tasks display an icon, click **Complete**.

Opening an Aborted Run

If you had accidentally aborted 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 aborted.

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 second rack to the private managed switch.

3 Follow "Image a Physical Rack," on page 32.

Viewing the VIA Log File

5

The log file displays information for all VIA services.

Procedure

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

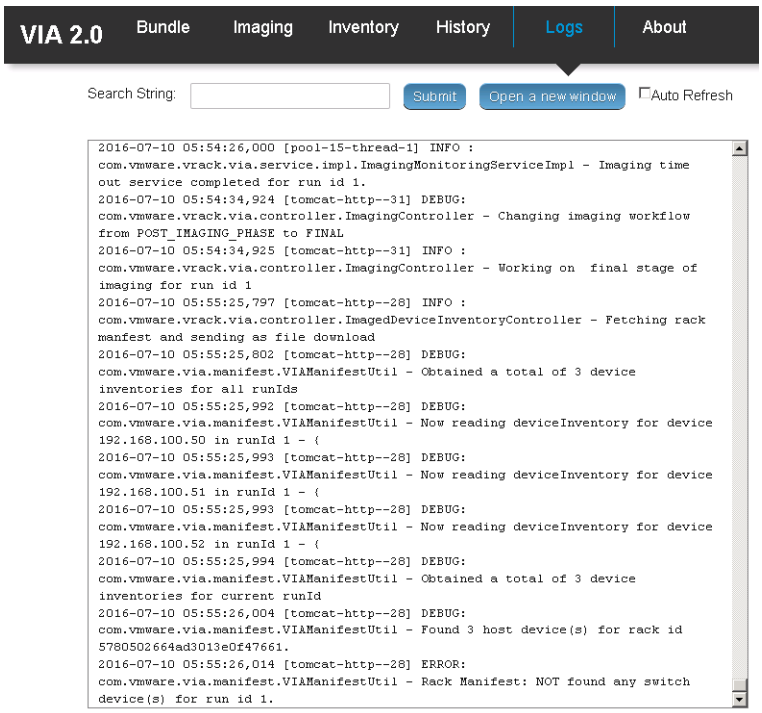

A consolidated log of VIA services is displayed sorted by the time stamp. A maximum of 500 entries is displayed at a time.

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 6

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

You can view the status of an imaging run by specifying its run ID. 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**.

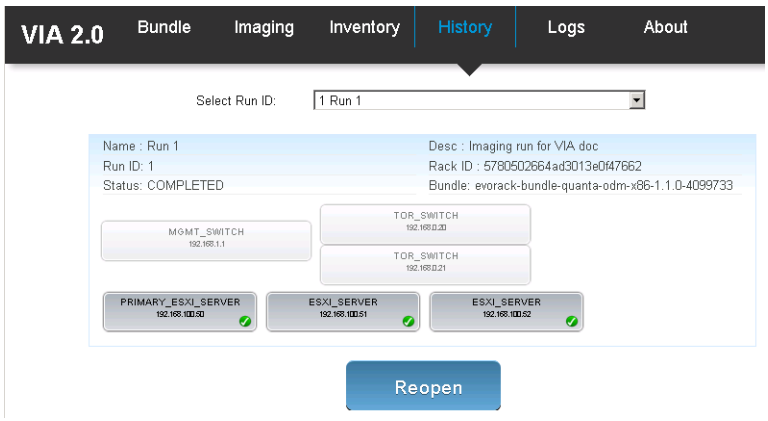

2 In **Select Run ID**, select the run ID for which you want to view the imaging history.

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

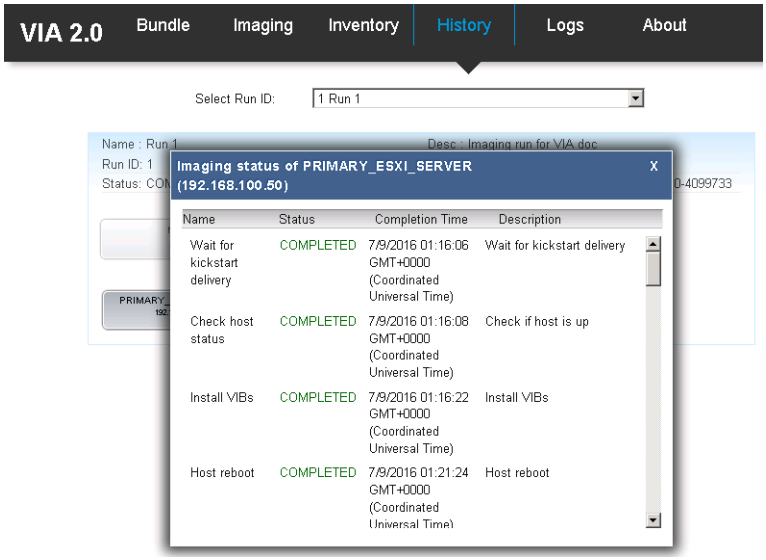

3 To view details for a device, click the expand icon next to the device.

4 To reopen a previous run, select the run ID and click **Reopen**.

You can continue imaging an aborted 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

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

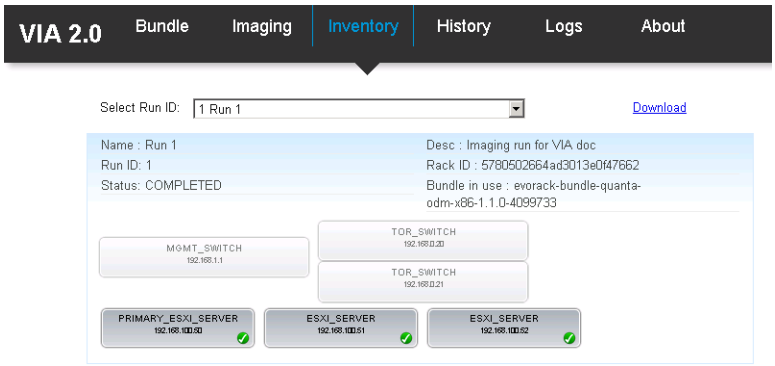

2 In Select Run ID, select run ID.

The device inventory for the selected imaging run is displayed.

3 To view details for a device, click the expand icon next to the device.

4 To download the rack inventory click **Download** and specify the directory where the file is to be saved.

The device inventory is saved as a JSON file.

BIOS Settings 7

The BIOS settings for each device in the physical rack must match the values given below.

This chapter includes the following topics:

Quanta Settings

Table 7-1. All components

Setting	Path to Setting
Set BIOS clock to current time	**Main > BIOS Information > System Time**

Table 7-2. Servers

Setting	Path to Setting	Value
Boot order	**System BIOS Settings > Boot Settings > Boot Sequence** 1 Use arrow key to reach the correct boot order number. 2 Press Enter. The boot devices are displayed. 3 Use arrow key to highlight Network. 4 Press Enter to select it.	Network first
HDD order	**System BIOS Settings > Boot > Fixed Boot Order priorities** 1 Use arrow key to reach the correct boot order number. 2 Press Enter. The boot devices are displayed. 3 Use arrow key to highlight SATADOM. 4 Press Enter to select it.	SATADOM first
Hyperthreading	1 Navigate to **System BIOS Settings > Advanced > Hyper-threading**. 2 Press Enter to enable.	Enabled
IPMI credentials		Default credentials
IPMI Network Settings		DHCP
Mode		Legacy

Table 7-2. Servers (Continued)

Setting	Path to Setting	Value
NUMA	1 Navigate to **System Bios Settings > Chipset > North Bridge > Numa** . 2 Press Enter to enable. Disabling node interleaving enables NUMA	Enabled
Power management	**System BIOS Settings > Advanced > CPU Power Management Configuration**	
EIST (P-states)		Enabled
Turbo Mode		Enabled
CPU C3 report		Disabled
CPU C6		Enabled
CPU Advanced PM Tunning / Energy Per BIAS		Balanced performance
PXE on 1G	1 Navigate to **System BIOS Settings > Advanced > Onboard Device Configuration** . 2 Select **Enabled Without PXE** for both the 1G NICs.	Disabled
PXE on 10G	10G NICs set by default to PXE. To verify, press Ctrl+s while the server is booting to enter the BIOS.	Enabled
VT	**System BIOS Settings > Processor Settings > Virtualization Technology Enabled**	Enabled

Dell Settings

Table 7-3. All components

Setting	Path to Setting
Set BIOS clock to current time	1 Navigate to **BIOS > System BIOS Settings > Miscellaneous Setting > System Time.**. 2 Click on the right panel to set time.

Table 7-4. Servers

Setting	Path to Setting	Value
Boot order	1 Navigate to **System BIOS > System BIOS Settings > Boot Settings > Bios Boot Settings** . 2 Click Boot Sequence. 3 Click the + icon to move Integrated NIC to the top. Use arrow key and + to move SD up to the top of the list	Network first
HDD order	1 Navigate to **System BIOS > System BIOS Settings > Boot Settings > Bios Boot Settings** . 2 Click Hard-Disk Drive Sequence . 3 Click the + icon to move the Internal SD card to the top.	SD Card first
Hyperthreading	**System BIOS > System BIOS Settings > Processor Settings > Logical Processor Enabled**	Enabled

Table 7-4. Servers (Continued)

Setting	Path to Setting	Value
IPMI credentials	1 To view the default IPMI credentials, navigate to **iDRAC SettingsUser Configuration**. 2 Do not change the default values for any settings: ■ User ID -> 2 ■ Enable User -> Enabled ■ User Name -> root ■ LAN User Privilege -> Administrator ■ Serial Port User Privilege -> Administrator ■ Change password -> blank **Note** VIA uses the default IPMI credentials which is root/calvin for Dell.	Default credentials
IPMI Network Settings	1 Navigate to **iDRAC Settings > Network > IPMI Settings > Enable IPMI Over LAN**. 2 Click **Enabled**.	Enabled on LAN
Mode		Legacy
NUMA	**System BIOS > System BIOS Settings > Memory Settings > Node Interleaving Disabled** Disabling node interleaving enables NUMA	Enabled
Power management	1 Navigate to **System BIOS > System Profile Settings > System > System BIOS Settings**. 2 Select Performance. This enables Turbo Boost.	
PXE on 1G Port 4	1 Navigate to **Device Settings > Integrated NIC 1 Port 3 Gigabit > NIC Configuration > Legacy Boot Protocol** . 2 Select **None**. 3 Repeat the above steps on the second integrated 1G NIC.	Disabled
PXE on 10G Port 2	1 Navigate to **Device Settings > Integrated NIC 1 Port 1 10G > NIC Configuration > Legacy Boot Protocol** . 2 Select **PXE**. 3 Repeat the above steps on the second integrated 10 G NIC.	Enabled
VT	**System BIOS > System BIOS Settings > Processor Settings > Virtualization Technology Enabled**	Enabled

Troubleshooting VIA 8

More information to be added to this chapter as EVO SDDC progresses through the Early Field Trials and RTP1.

This chapter includes the following topics:

Host failed to be imaged with error Unable to Establish IPMI v2 / RMCP+ Session

VIA was not able to power on a host and failed to image it.

Problem

After a host was powered off, VIA was unable to power it on. The following error was displayed.

Unable to establish IPMI v2 / RMCP+ session Unable to set Chassis Power Control to Up/On

Cause

VIA was unable to establish an IPMI v2 or RMCP+ session with the host.

Solution

1 Manually power on the host through DRAC.

2 On the **Imaging** tab, click the host that displayed the red icon and click Retry.

VIA continues imaging the rack.

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 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.

Cause

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

Solution

1 Connect to the management swtich 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-image-select -k`.

Index