

VxRail Deployment with Arista Networking Best Practices Guide

Converged Cloud Fabric

Table of contents

Introduction	3
VxRail Node Connectivity Options	4
Converged Cloud Fabric for VxRail Cluster	5
Bringing up Converged Cloud Fabric	6
CCF Infrastructure	6
Bringing up a VxRail Cluster	7
VxRail Deployment	9
Enabling vCenter Integration	9
Technical Resources	12
Summary	12

Introduction

Arista Converged Cloud Fabric (CCF) is an SDN-based leaf-spine (POD) fabric, which delivers consistency, automation, operational velocity, and visibility for enterprise customers. CCF tightly integrates with enterprise data center solutions, including hyper-converged Infrastructure (HCI), containers, and software-defined data centers.

The Dell EMC VxRail™ appliance is an HCI solution that consolidates compute and storage into a single, highly available, network-ready unit. VxRail has a simple, scale-out architecture, leveraging VMware vSphere® and VMware vSAN™ to provide server virtualization and software-defined storage.

Network connectivity is fundamental to the VxRail clustered architecture. Individual nodes act as a single system, providing scalability, resiliency and workload balancing by leveraging logical and physical networks.

A VxRail cluster consists of three or more VxRail nodes (ESXi hosts). Deployment, configuration, and management are handled by VxRail, allowing the compute capacity and the vSAN datastore to grow dynamically. VxRail Manager automatically discovers and configures each new node, and automatically adds it to the default vSphere Distributed Switch (VDS). vCenter propagates the port groups of the default VDS to the new node.

The following software versions are used to create this guide:

Solution	Product	Version
CCF for VxRail	Converged Cloud Fabric	CCF-6.0
	Dell Technologies VxRail	4.5

VxRail Node Connectivity Options

VxRail Appliances come in different models, offering various uplink connectivity options as shown in the diagram below. (Please check the Dell/EMC website for the latest model offerings).

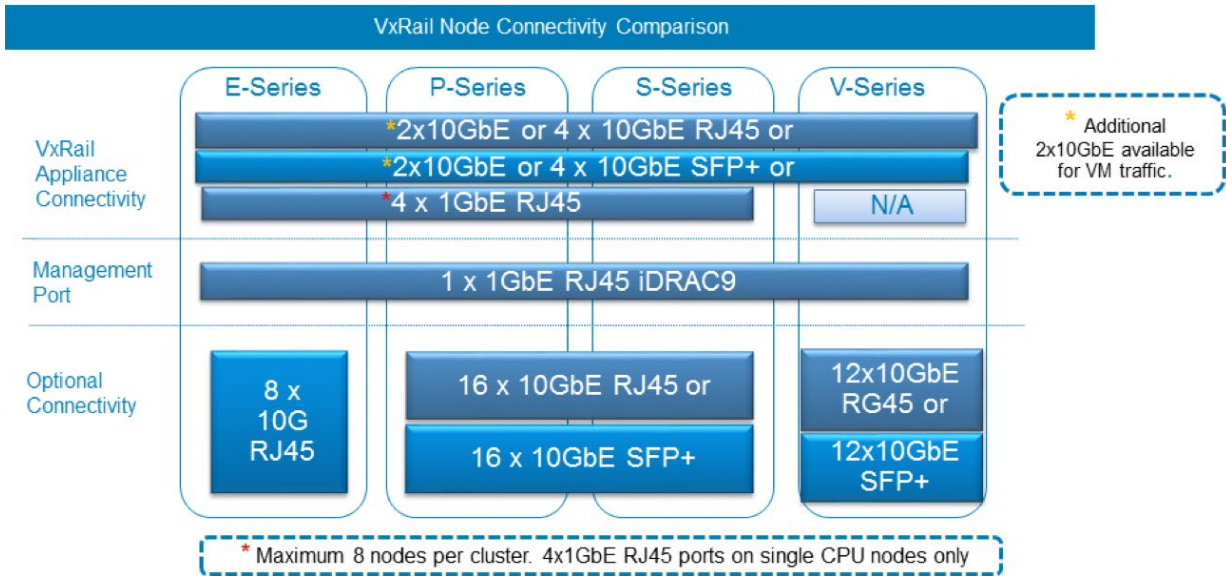


Figure 1: VxRail Connectivity Options

Converged Cloud Fabric for VxRail Cluster

Converged Cloud Fabric is an SDN-based leaf-spine networking fabric, which delivers consistency, automation, operational velocity, and visibility for enterprise customers. CCF tightly integrates with enterprise data center solutions, including hyper-converged Infrastructure (HCI), containers, and software-defined data centers.

When CCF integration with Dell EMC VxRail appliance is enabled, it provides following benefits (for more details see section “Enabling vCenter Integration”):

Leveraging this network-wide database - or NetDB - architecture, CloudVision focuses on three key pillars of functionality:

- Auto host detection and auto MLAG creation
- Auto L2 network creation
- End-to-end contextual visibility (physical + virtual)
- One-click troubleshooting

The following figure illustrates a three-node cluster connected to CCF leaf switches using 4 X 10GbE uplinks.

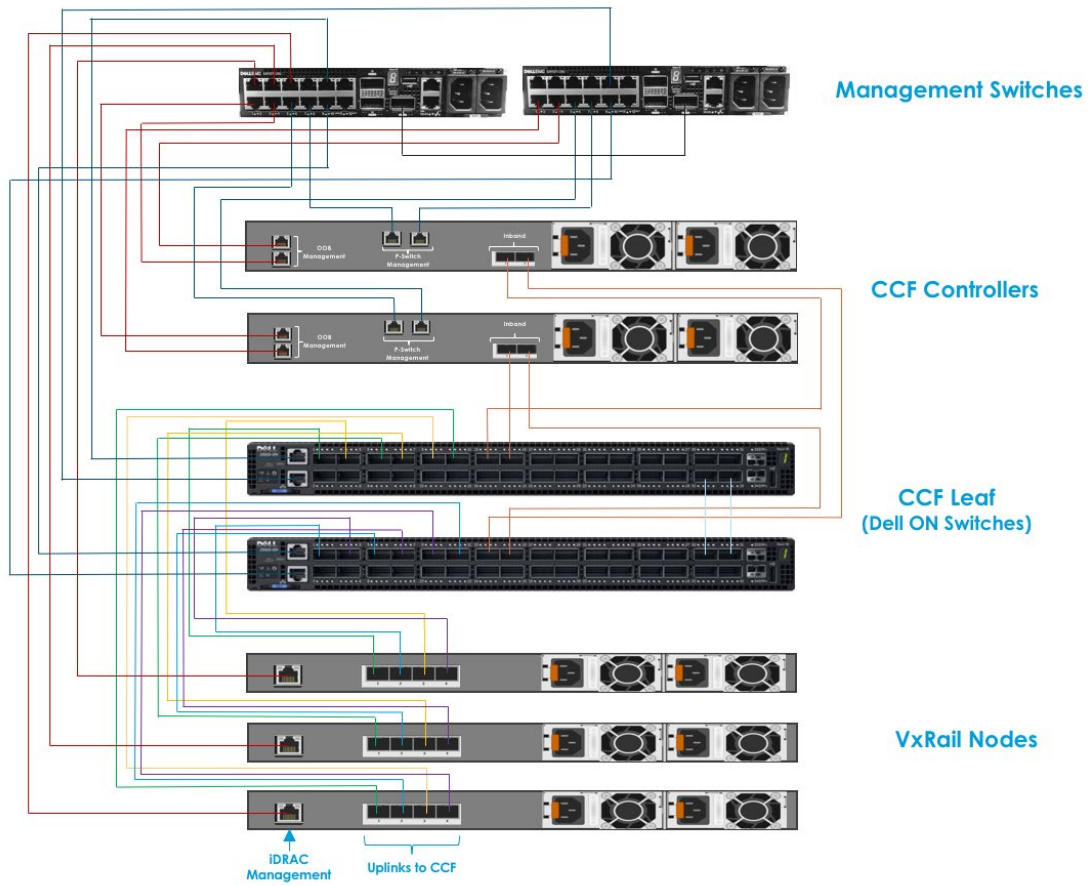


Figure 2: Three Node VxRail Deployment with CCF-EC

Node	Uplink 1	Uplink 2	Uplink 3	Uplink 4
Node1	TOR1-Eth1	TOR2-Eth1	TOR1-Eth2	TOR1-Eth2
Node2	TOR1-Eth3	TOR2-Eth3	TOR1-Eth4	TOR1-Eth4
Node3	TOR1-Eth5	TOR2-Eth5	TOR1-Eth6	TOR1-Eth6

Bringing up Converged Cloud Fabric

To bring up Converged Cloud Fabric, please refer to deployment steps in Chapter 1-4 in the Converged Cloud Fabric Deployment Guide. While bringing up the fabric, please keep the best practices listed below in mind.

Category	Description	Must / Recommended
CCF Infrastructure	Assign dedicated VLAN for CCF management network	Recommended
	Allow IPv6 traffic on p-switch management network	Must
	Use dual management switches or stack switches for p-switch management connections	Recommended
	CCF Controller pair should be placed in separate racks (in case of rack-to-PDU dependency)	Recommended
	Inter-leaf-switch connection (peer) should be at least two links for redundancy	Recommended
	Use Hardware Compatibility List (HCL) supported optics and cables	Must
	Dual connected host (e.g ESXi server), services (e.g FW, LB), etc. in same rack (same leaf-group)	Recommended
	mDNS must be allowed to be flooded over the p-switch management network	Must
	MTU on the p-switch management network must be configured to 9K	Must
VMware	All ESXi hosts connected to Converged Cloud Fabric switches must have unique hostnames. ("Domain" field should not be empty. It is required to use fully-qualified-domain names (FQDN) for all ESXi hosts)	Must
	When not using LACP, set the network teaming failback delay on ESXi hosts (Net.TeamPolicyUpDelay) to 30 seconds. (to provide better traffic failover when leaf switches reboot)	Must
	vSphere inventory property naming should not have special characters including spaces. (For example vSwitch/DvSwitch names and vSwitch/DvSwitch portgroup names must begin with lower or upper-case alphabetic characters and can contain upper or lower-case alphanumeric characters or a dash, or asterisk, or a period [.])	Must
	To properly form a Converged Cloud Fabric interface-group ESXi hostname+domain name and vSwitch/DvSwitch name combination should be under 255 characters	Must
	CCF automation for ESXi workloads, make sure no manual CCF interface groups (LAG) are configured on the CCF controller connected to ESX hosts	Must
	Use LLDP as the discovery protocol instead of CDP	Recommended

Bringing up a VxRail Cluster

VxRail infrastructure carries traffic in the following categories:

- In-Band Management
- vMotion
- vSAN
- VM Traffic

In-band management traffic includes all VxRail, vCenter Server, and ESXi communication. The management VLAN also carries traffic for (1) vRealize Log Insight and (2) VMware Loudmouth required for VxRail cluster initialization (Loudmouth is a service that discovers and configures nodes on the network during Initial configuration and appliance expansion). All management traffic can be untagged or you can specify a VLAN tag.

Note: Connectivity between the vMotion VMkernel interfaces and the vSAN VMkernel interfaces is required when bringing up the VxRail cluster. This is achieved via the segment configuration covered in Step 2 below.

For L2/L3 isolation, CCF uses a construct called Enterprise Virtual Private Cloud or E-VPC. The current deployment example uses the CCF VxRail E-VPC (tenant) to provide in-band connectivity for management traffic, for vSAN and vMotion traffic, and for bootstrapping the VxRail cluster.

To deploy the VxRail solution, complete the following steps.

Procedure

Step 1: Collect the necessary information about the management and external network

Information to collect	Example
Management VLAN for ESXi, vCenter and VxRail communication	VLAN 4001
VLAN for vMotion	VLAN 4002
VLAN for vSAN	VLAN 4003
VLAN for Workload	VLAN 4004

Step 2: Create the CCF VxRail E-VPC (tenant).

Manually configure the following E-VPC, segment, and logical router on the CCF active controller.

```
evpc-tenant VxRail
  segment VxRail-Management
    member interface-group any vlan 4001
    member switch any interface any vlan 4001
  segment VxRail-vMotion
    member interface-group any vlan 4002
    member switch any interface any vlan 4002
  segment VxRail-vSAN
    member interface-group any vlan 4003
    member switch any interface any vlan 4003
    segment VxRail-workload
      member interface-group any vlan 4004
      member switch any interface any vlan 4004
```

The VxRail-Management segment defined within the VxRail E-VPC (tenant) carries the management traffic for vCenter and ESXi hosts. The following segment membership rule assigns all traffic from VxRail nodes with VLAN 4001 to the VxRail-Management segment.

```
segment VxRail-Management
  member interface-group any vlan 4001
  member switch any interface any vlan 4001
```

The following segment membership rule assigns traffic with VLAN 4002 to the VxRail-vMotion segment.

```
segment VxRail-vMotion
  member interface-group any vlan 4002
  member switch any interface any vlan 4002
```

The following segment membership rule assigns traffic on VLAN 4003 to the VxRail-vSAN segment.

```
segment VxRail-vSAN
  member interface-group any vlan 4003
  member switch any interface any vlan 4003
```

The following segment membership rule assigns traffic on VLAN 4004 to the VxRail-workload segment.

```
segment VxRail-vSAN
  member interface-group any vlan 4004
  member switch any interface any vlan 4004
```

Step 3: (Optional) External connectivity

When a workstation connecting to the VxRail UI using a web browser is not connected directly to the CCF leaf switches, the connection must be routed by the router external to CCF to establish connectivity to the VxRail nodes.

Connectivity between the VxRail E-VPC and the external network is established, as shown in the following figure.

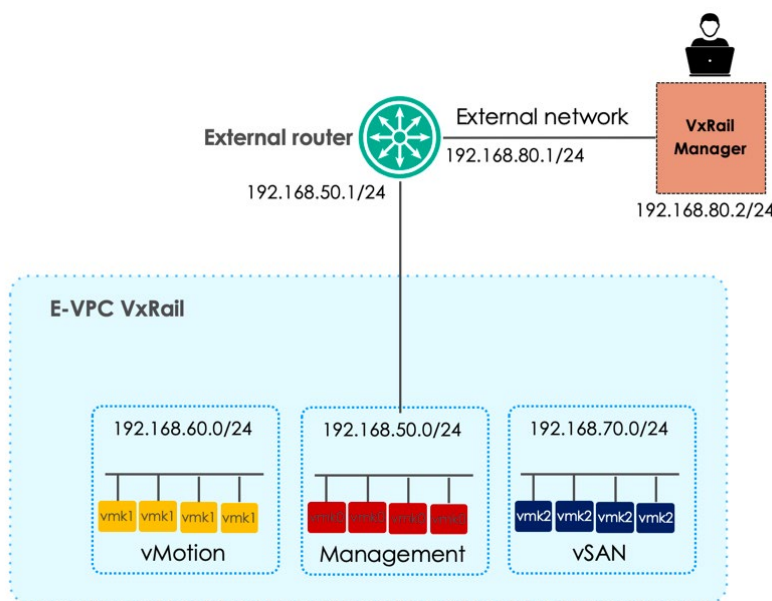


Figure 5: Establishing connectivity between VxRail E-VPC and external networks

Here we are assuming that the interface-group connecting to the external router will use VLAN 4001 and hence no explicit membership rule is not required under the VxRail-Management segment.

Layer 3 routing between the management segment on CCF and the external network is taken care of at the external router.

VxRail Deployment

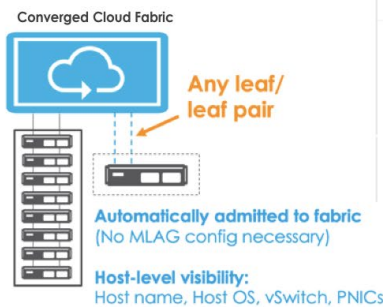
At this point, the network is ready for VxRail deployment. Please follow the [Dell/EMC VxRail deployment steps](#) (Chapter 9) to form the VxRail Cluster.

Enabling vCenter Integration

CCF vCenter integration will provide the following benefits

1. Automatically detect where ESXi nodes connect to the fabric and create interface-group for the node. This makes Node addition/removal on VxRail cluster zero touch for the network

Add VxRail Node...



Name	Description	Leaf Group	State	Admin Status	Mode	Auto-Discovered	Backup Mode	Filter	Backu	Prin	Aut
ESX6-111.tme.bigswitch.com-DSwitch-HQ-Cluster1	Interface group for virtual switch DSwitch-HQ-Cluster1 in ESXi host ESX6-111.tme.bigswitch.com	R1	Up	Up	LLDP	-	Static				

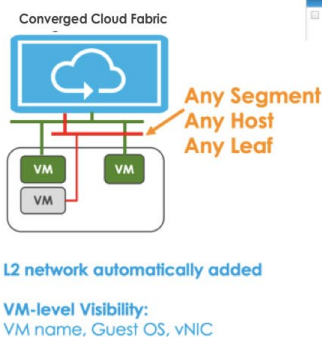
Switch	Switch MAC	Interface Name	Description	Status	Spine Switch	Leaf Switch	Virtual Switch	Operational	Physical
R1L1	cc:37:ab:2c:98:34	ethernet21		Up	-	-	-	-	-
R1L2	78:72:cf:1e:f1:84	ethernet21		Up	-	✓	-	-	-

Host Name	Interface Name
ESX6-111.tme.bigswitch.com	vminic4
ESX6-111.tme.bigswitch.com	vminic5

As soon as the ESXi nodes are physically connected to CCF, they get auto discovered and provisioned as per the teaming policy of the VDS uplinks from each node, irrespective of the number of hosts connected to the CCF. Using Converged Cloud Fabric, there is no need to manually configure the switch and interface where the host connects, thus simplifying Host Network provisioning. No hard-wired port mapping needed -- a server link can be connected to any speed-appropriate switch port. CCF automatically re-provisions for the new port. Also any server can be placed in or moved to any rack at any time -- CCF controller does the heavy lifting of automatic logical-to-physical mapping through SDN intelligence while providing full topology visibility to the network admin.

2. Automate the segment (L2 Network/VLAN) creation and removal on the CCF for the user configured port groups

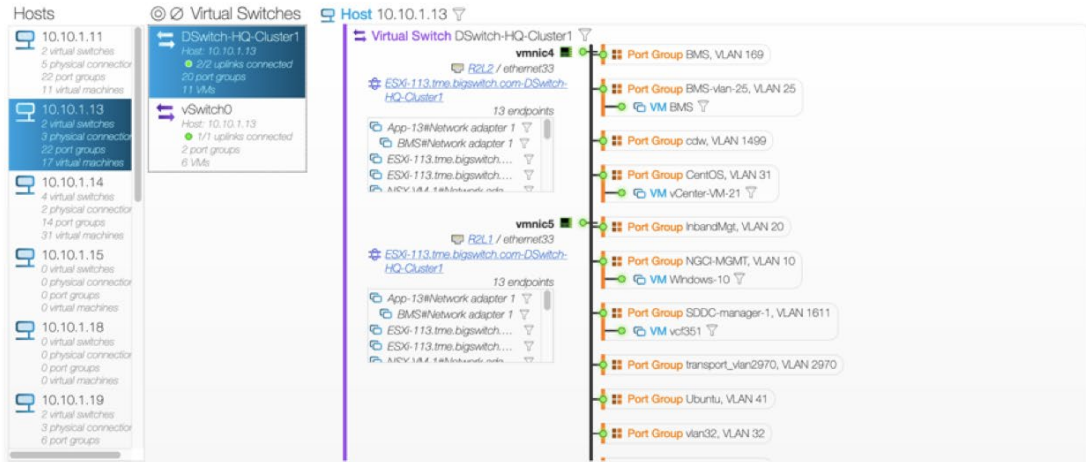
Add ESXi VM...



Tenant	Segment	MAC	Vendor	Name	Source	Shutdown Status	State	Description
YC50	YC50-30	98:58:56:0a:e7:f8	VMware, Inc.	vm-75-00-50-56-a8-e7-f8	Config	Up	Active	CentOS-15m-75 [Network adapter 1] on ESXi Host: 10.2.18.2 [VMware Tools NOT correctly installed/toolsNotInstalled]

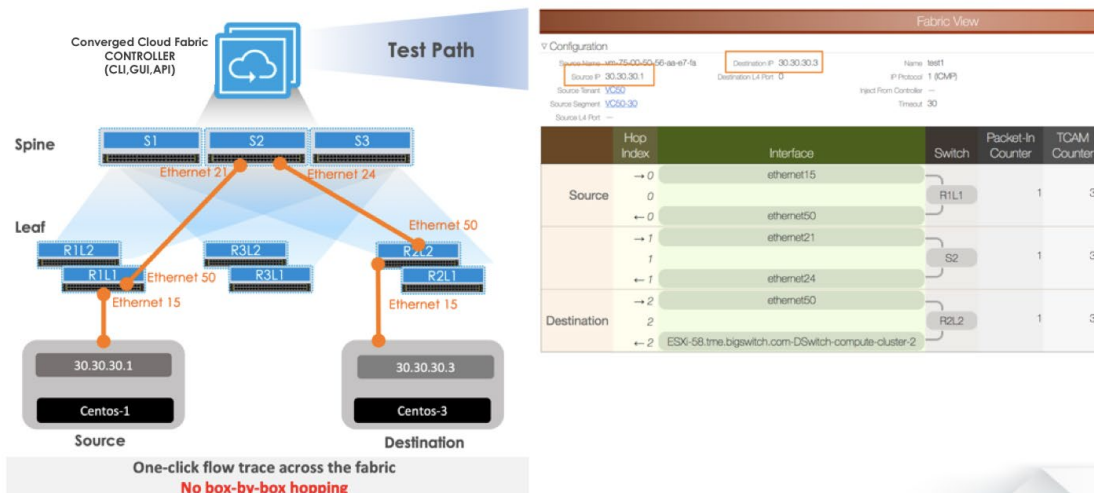
CCF creates E-VPC for vCenter, allowing logical isolation and delegated administration, to automate transport VLAN/VLANs provisioning within the E-VPC, and trunk the VLANs on the appropriate host interfaces. Admins need not perform any manual configuration box-by-box. As and when more networks are created on vCenter, CCF automatically adds the configuration thus reducing wait time to provisioning a new host. Even when the host is moved from one rack to another, no network provisioning is required.

3. Provide VM, Host and VDS level visibility to the Network Admin from CCF GUI



CCF provides visibility into the vCenter networking in a single dashboard, making it easy for network & virtualization admins to get an end-to-end picture.

4. Provide VM-to-VM and vmk-to-vmk hop by hop packet tracing using Fabric Trace.



With CCF Fabric Trace, admins can trace end-to-end packets between any VMs or vmk interfaces across the fabric, with just one click and get hop by hop packet stats, thus enabling admins to restore services much more rapidly.

To enable vCenter integration, complete the following configuration. Exclude the Management, vSAN and vMotion VLANs as shown in the following example running-config.

Excluding Management, vSAN and vMotion VLANS in CCF vCenter configuration

```
! vcenter
vcenter VC65
  automation-level full
  exclude vlan 4001-4003
  hashed-password 5606520c351828584d014b5546474e
  host-name 10.10.1.2
  manage-segment-for-vlan-range start-vlan untagged end-vlan 4095 tenant VC65
  user-name administrator@vsphere.local
```

Once the vCenter integration is enabled, all the Day1/Day2 operations like host addition/removal, portgroup addition/removal, vMotion events are automated

Technical Resources

Converged Cloud Fabric: <https://www.bigswitch.com/products/big-cloud-fabric>

CCF + VxRail video demo: <https://www.youtube.com/watch?v=A3d9ksYxX8o>

CCF on-line labs: <https://CCF-labs.arista.com/>

Summary

With Arista Converged Cloud Fabric solutions for VxRail, Arista Networks offers the most comprehensive, flexible and highly automated solution in the industry. This solution has been tested and validated and is supported by Arista Networks. Customers who want to realize the operational benefits of a hyper-converged solution from Dell EMC can deploy it with Arista's best-in-class data center networking solutions.

Release Date: March, 2020; Version: 1.0

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office

9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office 1390

Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office

9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062



Copyright © 2020 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. May 4, 2020