ARISTA

White Paper

Converged Cloud Fabric for VMware NSX-T

Document Version 1.0

Introduction

Organizations are constantly striving to simplify their operational environment to support dynamic business priorities. As they embark on this journey of digital transformation to modernize their data centers, they are rapidly embracing cloud-style principles like self-service operations, agility, scale-out architecture, resiliency and unified management to name a few.

By leveraging VMware Software-Defined Data Center (SDDC) technologies, organizations are able to drive agility and cost efficiencies for their application workloads. This holistic software-defined approach encompasses automated application deployment across both physical and virtual infrastructure. NSX-T is the networking and security component of the VMware software-defined data center (SDDC) stack. Using NSX-T manager, admin can virtualize the networking components and build logical overlay network topologies with virtual switches, virtual routers, virtual firewalls and virtual load balancers on-demand. This makes the overlay network as agile as the application deployment.

Legacy underlay networks, however, still prove to be a bottleneck as they have been traditionally challenging for data center administrators to design and configure. It is imperative to provide cloud-style agility, experience and operational simplicity for underlay networks as well. Network needs to be able to provide seamless Day0/Day1/Day2 operations in a heterogeneous application environment with the flexibility of deploying any virtualization, container and Hyper-Converged Infrastructures (HCI).

Automation across physical and virtual networks therefore becomes a critical aspect of the SDDC-automated infrastructure. In addition, gaining visibility across physical and virtual networks is becoming paramount for network and VMware administrators, as troubleshooting has been challenging with traditional networks.

Solution

Arista Converged Cloud Fabric (CCF) approach simplifies data center networking by providing operational consistency, visibility and governance. The platform embraces the same design principles adopted by public cloud providers to offer a fully cloud-like experience for on- premise data centers. Arista CCF empowers organizations to transform and future-proof their networks and meet the scale and performance requirements of the new digital economy.

The CCF controller operates as "one logical switch", which removes complexity and automates Day0/Day1/Day2 operations, delivers Network-as-a- Service through cloud-style Enterprise Virtual Private Cloud (E-VPC), and provides contextual analytics for deep visibility.

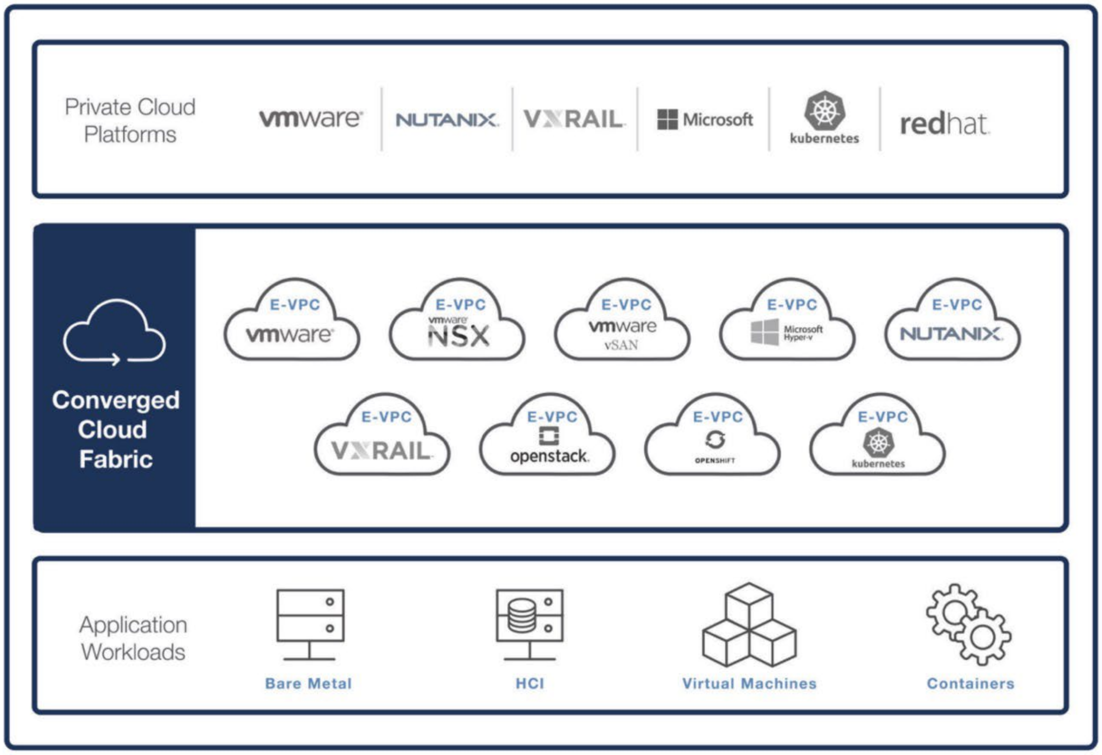

Figure 1: E-VPCs on CCF

Challenges when deploying VMware NSX-T on a legacy box-by-box underlay

In the following section we will discuss some of the challenges and inefficiencies when the underlay network is a legacy box-by-box network and how CCF mitigates these challenges and inefficiencies.

1. **Host network provisioning:** ESXi/KVM nodes need to be connected to the physical fabric, and those specific ports on TOR switches need to be configured manually with appropriate configurations for LAGs/LACP, etc. depending on the teaming policy of the N-VDS uplinks from each node. Each rack can have 20 to 40 ESXi/KVM hosts, resulting in 40-80 interfaces/LAGs/LACP configurations, which significantly increases the time to service enablement and also the scope of error.
2. **Transport VLAN configuration:** In order for the GENEVE tunnels to be established between hosts, the underlay network needs to provide connectivity between Tunnel Endpoints (TEP) defined on these hosts. Network admins would need to define and trunk transport VLANs on multiple switches, increasing the number of touch points and possibility of misconfigurations. In case of VLAN-based (non-GENEVE) logical switches, the number of VLANs that need to be defined and trunked in the underlay network would increase proportionally, leading to increased setup time and scope of misconfigurations/errors.
3. **Visibility:** In order to get end-to-end visibility, network & virtualization admins would need to gather overlay and underlay connectivity information from multiple consoles, and perform underlay/overlay correlation manually, which is extremely time consuming. Getting historical information to isolate a problem is a challenging and cumbersome process as it needs log scraping on multiple switches. As one might infer, the problem gets worse as the number of switches in the underlay increases.
4. **Troubleshooting:** To troubleshoot any underlay connectivity issues, network admins would need to perform box-by-box hopping to figure out the end-to-end path of the packets and isolate the switch causing the issue. This can significantly increase the time to restore services, leading to bad customer experience and potential revenue loss.

As highlighted above, a legacy box-by-box underlay can be cumbersome, error prone, time consuming and inefficient – it just does not scale operationally. You need an underlay that can operate at the speed of the VMs/Containers.

Deploying VMware NSX-T on CCF underlay

Let's discuss how CCF makes the life of a network admin easy by making the underlay more agile. As soon as the integration is enabled between NSX-T manager and CCF, an E-VPC is created for the NSX-T manager.

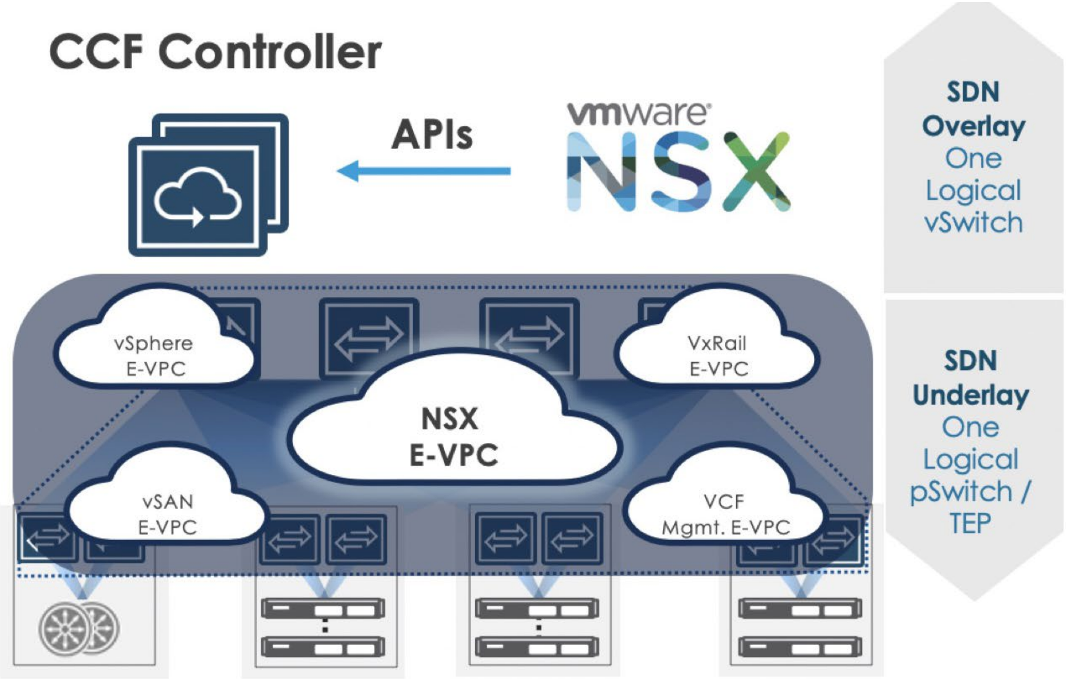

Figure 2: VMware NSX-T deployment on CCF

Host Network provisioning

As soon as the ESXi/KVM nodes are physically connected to CCF, they get auto discovered and provisioned as per the teaming policy of the N-VDS/VDS uplinks from each node, irrespective of the number of hosts connected to the CCF underlay. 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.

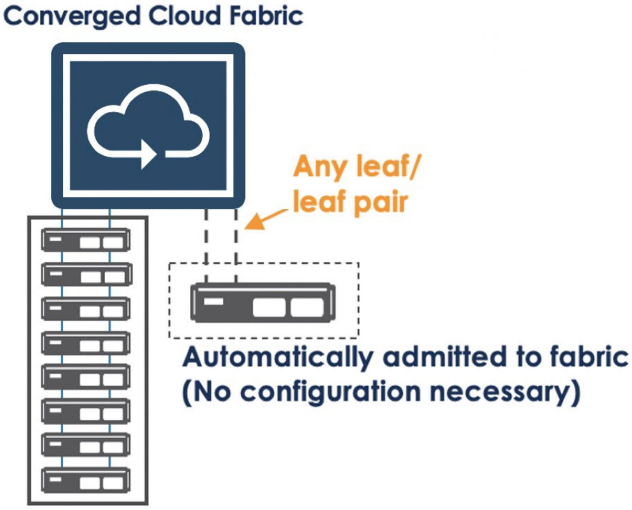
Converged Cloud Fabric
Any leaf/ leaf pair
Automatically admitted to fabric (No configuration necessary)

Figure 3: CCF Host automation

Transport VLAN provisioning

Just like the public clouds use virtual private cloud (VPC) logical construct to build multi-tenant L2/L3 networks, CCF leverages an AWS-style VPC on-prem construct -- called Enterprise VPC (E-VPC) to deliver Cloud-Network-as-a-Service operational experience.

CCF creates E-VPC for NSX-T, 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 we saw with the legacy approach.

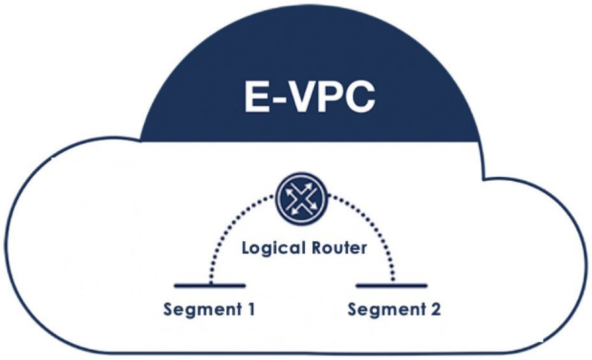

Figure 3: CCF Host automation

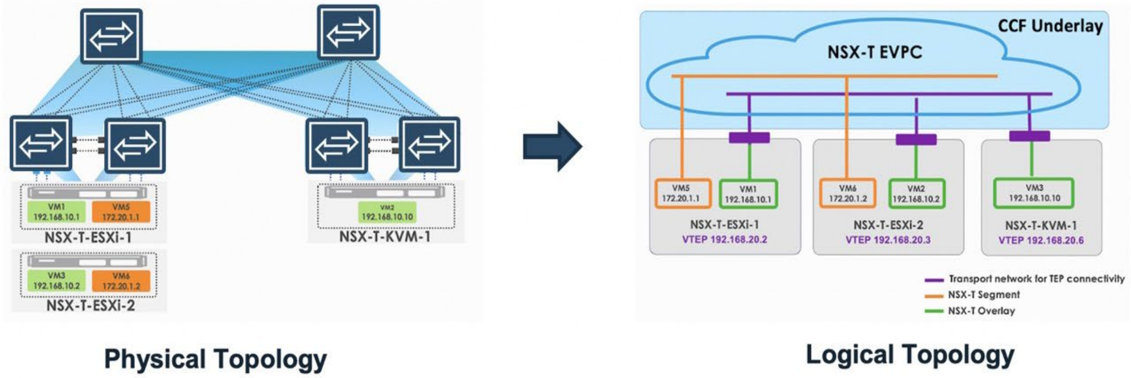

Figure 5: VMware NSX-T deployment with CCF

As and when more networks are created on NSX-T, 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.

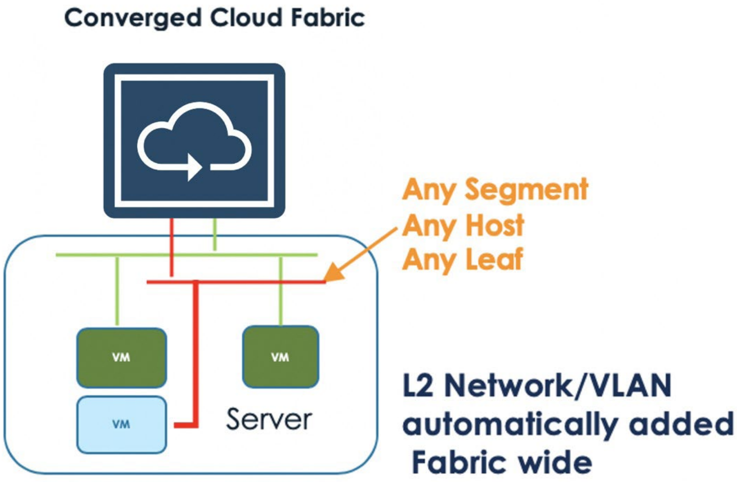

Figure 6: CCF Network Automation

Visibility

CCF provides visibility into the NSX networking environment as well as underlay networking in a single dashboard, making it easy for network & virtualization admins to make overlay/underlay correlation and get an end-to-end picture.

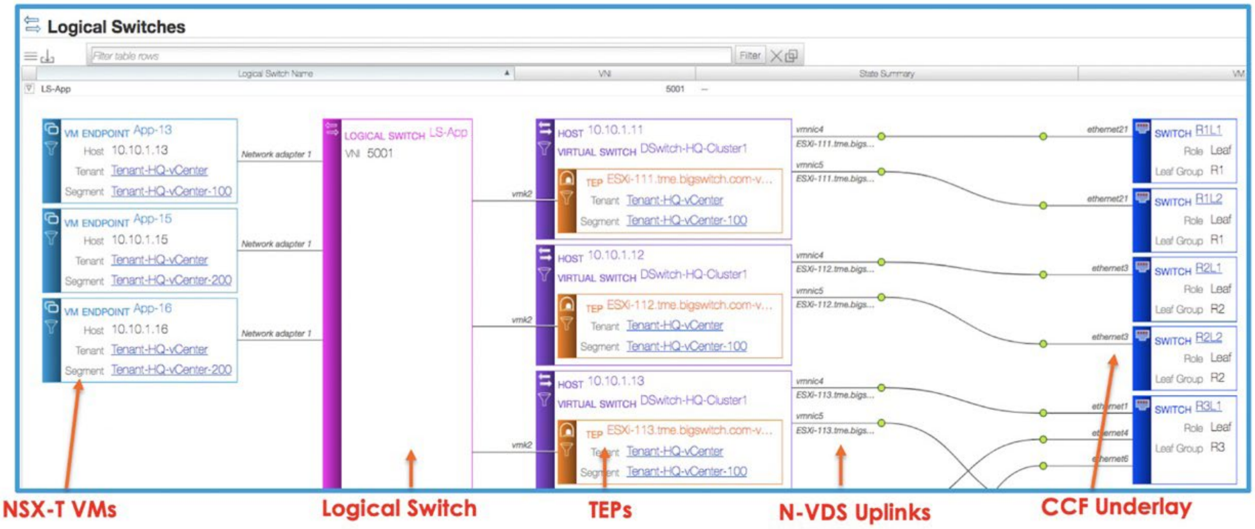

Logical Switches

Filter table rows | Filter

Logical Switch Name	VNI	State Summary	VM
LS-App	5001	–	

VM ENDPOINT App-13
Host 10.10.1.13
Tenant Tenant-HQ-vCenter
Segment Tenant-HQ-vCenter-100

VM ENDPOINT App-15
Host 10.10.1.15
Tenant Tenant-HQ-vCenter
Segment Tenant-HQ-vCenter-200

VM ENDPOINT App-16
Host 10.10.1.16
Tenant Tenant-HQ-vCenter
Segment Tenant-HQ-vCenter-200

LOGICAL SWITCH LS-App
VNI 5001

HOST 10.10.1.11
VIRTUAL SWITCH DSwitch-HQ-Cluster1
TEP ESXi-111.tme.bigswitch.com-v...
Tenant Tenant-HQ-vCenter
Segment Tenant-HQ-vCenter-100

HOST 10.10.1.12
VIRTUAL SWITCH DSwitch-HQ-Cluster1
TEP ESXi-112.tme.bigswitch.com-v...
Tenant Tenant-HQ-vCenter
Segment Tenant-HQ-vCenter-100

HOST 10.10.1.13
VIRTUAL SWITCH DSwitch-HQ-Cluster1
TEP ESXi-113.tme.bigswitch.com-v...
Tenant Tenant-HQ-vCenter
Segment Tenant-HQ-vCenter-100

SWITCH R1L1
Role Leaf
Leaf Group R1

SWITCH R1L2
Role Leaf
Leaf Group R1

SWITCH R2L1
Role Leaf
Leaf Group R2

SWITCH R2L2
Role Leaf
Leaf Group R2

SWITCH R3L1
Role Leaf
Leaf Group R3

Figure 7: Underlay/Overlay correlation

With CCF Fabric analytics, network admins can not only see events, errors, logs, performance stats from all the underlay switches in the fabric, but also all the events from NSX manager and vCenter, all in a single console. Network admins can now easily visualize the current state of the physical (underlay) and virtual (overlay) network or go back in time and perform historical analysis right from a single dashboard.

Troubleshooting

With CCF Fabric Trace, admins can trace end-to-end packets between any connected TEPs across the fabric, with just one click and get hop by hop packet stats, thus enabling admins to restore services much more rapidly as compared to box-by-box underlay networks.

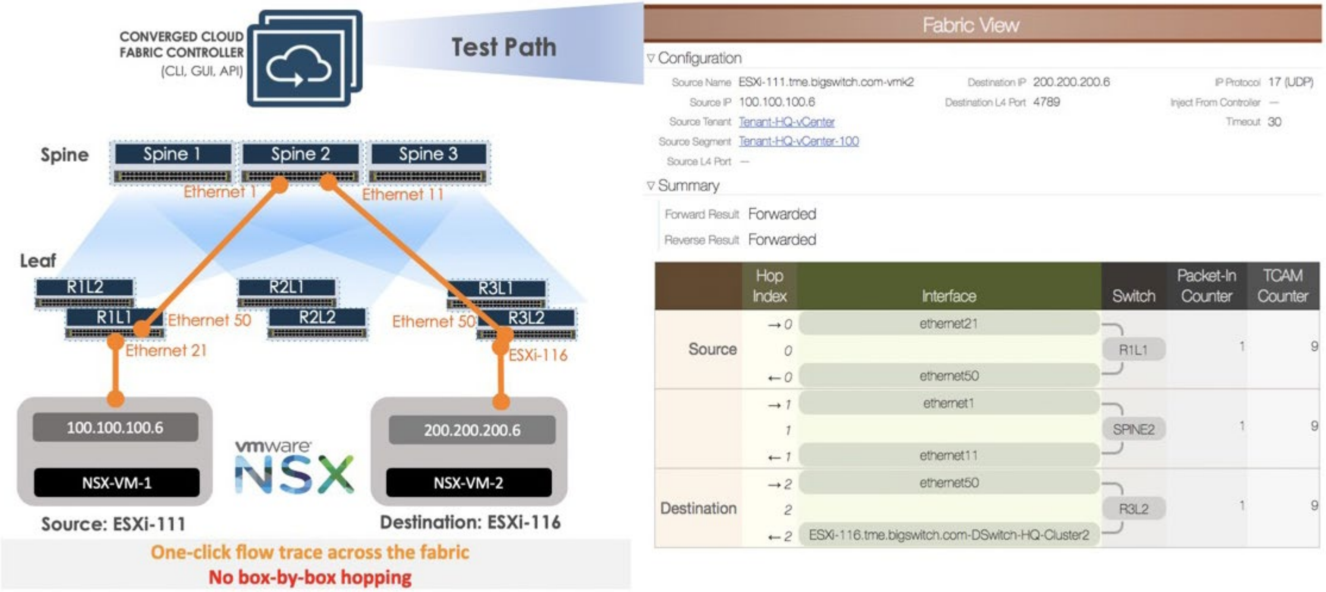

Fabric View

▽ Configuration

Source Name ESXi-111.tme.bigswitch.com-vmk2
Source IP 100.100.100.6
Source Tenant Tenant-HQ-vCenter
Source Segment Tenant-HQ-vCenter-100
Source L4 Port —
Destination IP 200.200.200.6
Destination L4 Port 4789
IP Protocol 17 (UDP)
Inject From Controller —
Timeout 30

▽ Summary

Forward Result Forwarded
Reverse Result Forwarded

	Hop Index	Interface	Switch	Packet-In Counter	TCAM Counter
Source	→ 0	ethernet21	R1L1	1	9
	0				
	← 0	ethernet50			
	→ 1	ethernet1	SPINE2	1	9
	1				
	← 1	ethernet11			
Destination	→ 2	ethernet50	R3L2	1	9
	2				
	← 2	ESXi-116.tme.bigswitch.com-DSwitch-HQ-Cluster2			

Figure 8: One click troubleshooting with CCF Fabric Trace

Summary

With Arista CCF, with its underlay orchestration capabilities, enhanced end-to-end visibility and one-click troubleshooting, provides a perfect underlay for VMware NSX-T

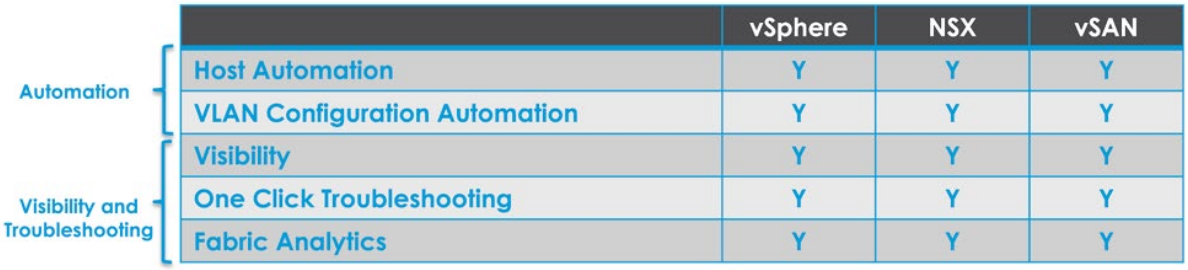

		vSphere	NSX	vSAN
Automation	Host Automation	Y	Y	Y
	VLAN Configuration Automation	Y	Y	Y
Visibility and Troubleshooting	Visibility	Y	Y	Y
	One Click Troubleshooting	Y	Y	Y
	Fabric Analytics	Y	Y	Y

Figure 9: CCF Integration for VMware SDDC suite

Technical Resources

Converged Cloud Fabric: https://www.arista.com/en/products/converged-cloud-fabric

Hands on labs: http://ccf-labs.arista.com

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