Converged Cloud Fabric for VMware vSphere

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.

Networking, however, has traditionally been challenging for data center administrators to design and configure. It is imperative to provide cloud-style agility, experience and operational simplicity for on-premises data center networking. 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 cloud-like experience for on- premise data centers. Arista CCF empowers organizations to transform and future-proof their networks for the scale and performance requirements of the new digital economy.

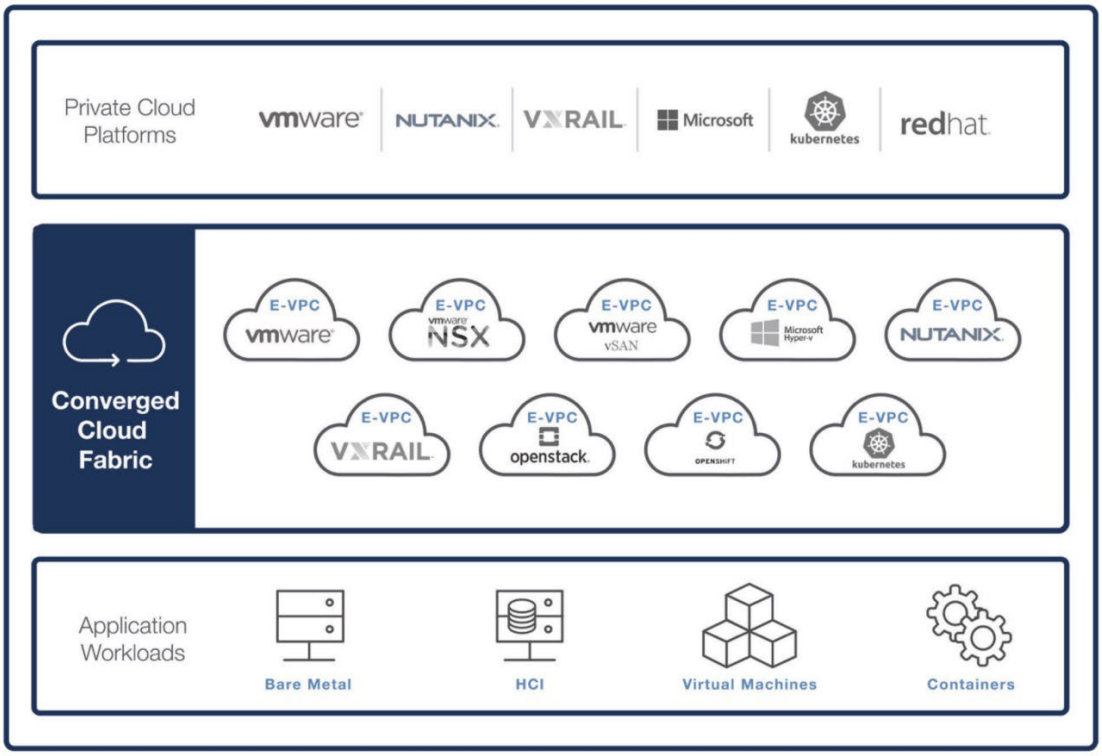

Figure 1: E-VPCs on CCF

The CCF controller operates as "one logical switch", which removes complexity and automates Day0/Day1/Day2 operations.

CCF empowers application teams with self-service experience for on-prem data centers as offered in public cloud.

CCF's Enterprise VPCs (E-VPC) allows rapid VM and container deployments using a logical network abstraction. E-VPC simplifies on-prem networking by expressing networking rules within the E-VPC and alleviates provisioning dependency of network hardware constructs (ports, VLANs, VRFs).

By transitioning to a logical network architecture built around externalized software control, the enterprises can accomplish cloud-style self-service networking -- a critical piece of the digital transformation journey .

In this document we will focus on how CCF's integration with the VMware vSphere server virtualization environment helps to automate Day0/Day1/Day2 operations.

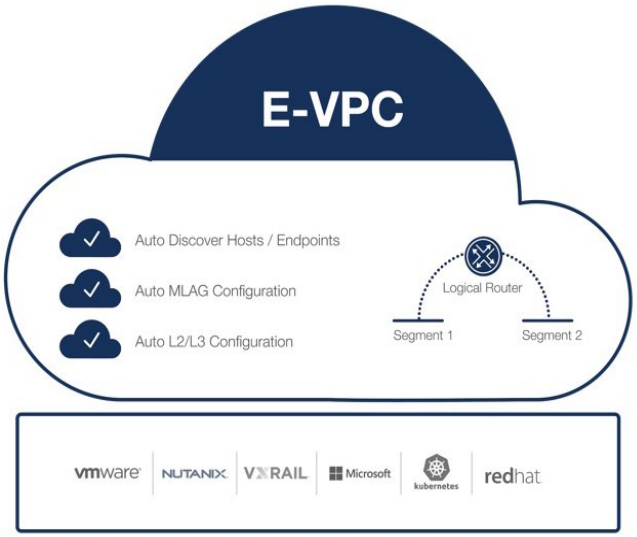

Challenges when connecting VMware vSphere workloads using a legacy box-by-box network

1. **Host network provisioning:** ESXi 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 VDS uplinks from each node. Each rack can have 20 to 40 ESXi hosts, resulting in 40-80 interfaces/LAGs/LACP configurations, which significantly increases the time to service enablement and also the scope of error.
2. **L2 Network VLAN configuration:** In order to establish communication between hosts on different hosts, the network needs to provide connectivity between these hosts. Network admins would need to define and trunk VLANs on multiple switches, increasing the number of touch points and possibility of misconfigurations. As the number of VLANs that need to be defined and trunked in the network increases, it leads to increased setup time and scope of misconfigurations/errors.
3. **Dynamic policy migration:** As virtual endpoints move between hosts that connect to different TOR switches, network admin needs to configure the appropriate VLANs on the new TOR and trunk them on the link connecting to the host. Also, the VLANs would need to be pruned from the original TOR if there are no more endpoints left in the VLAN. To get around this, one can preconfige all the VLANs on all the TORs and trunk them towards all the connected hosts, which helps with seamless connectivity for the endpoint mobility but also presents security challenges.
4. **Visibility:** In order to get end-to-end visibility, network & virtualization admins would need to gather information about the physical network and the virtual environment by going to multiple consoles, 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 network increases.
5. **Troubleshooting:** To troubleshoot any network 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 network can be cumbersome, error prone, time consuming and inefficient – it just does not scale operationally. You need a network that can operate at the speed of the VMs/Containers.

Deploying vSphere with CCF

The CCF Controller acts as a single pane of glass for fabric configuration and integrates with the vCenter for L2 physical network automation. As new VMs are created in vCenter, they are automatically learned in the fabric, and network policies are auto-migrated upon VMs migrating using vMotion.

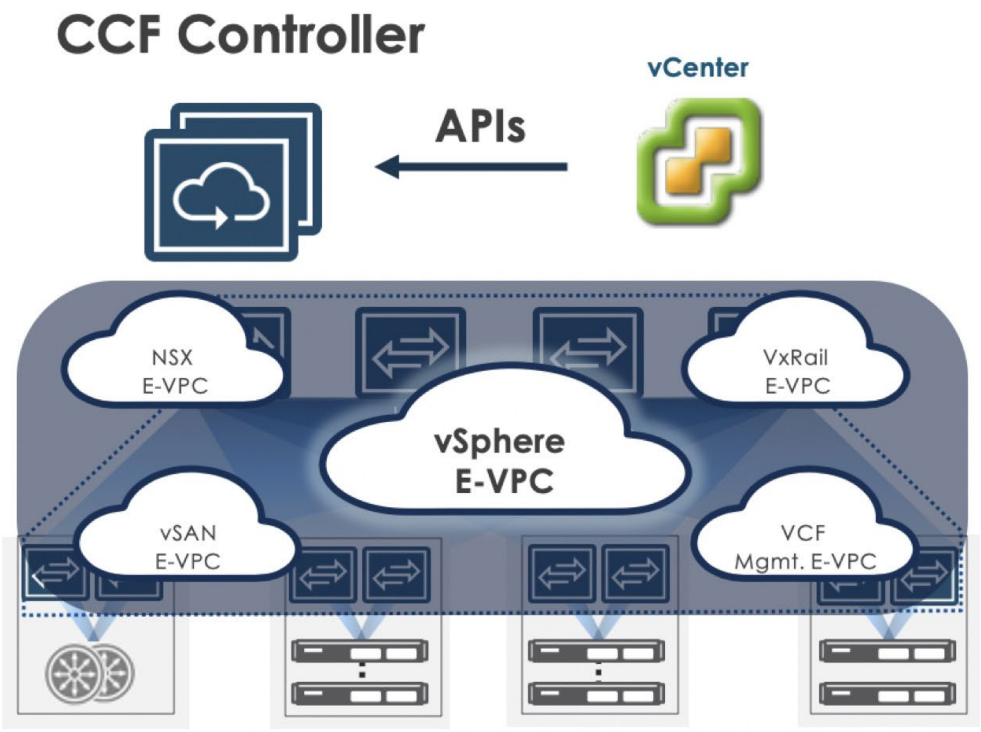

Figure 2: CCF vSphere integration

The CCF Controller also acts as a single pane of glass to provide VM-level, as well as vMotion, visibility across the physical fabric. This visibility coupled with advanced analytics enables fabric-wide troubleshooting—offering operational simplicity compared to legacy approaches. Net Admins gain full visibility of server virtualization environments, helping rapid resolution of issues.

Network Automation with vSphere and CCF

The CCF/vSphere solution streamlines application deployment workflows by automating the physical network configuration for VMware virtual workloads. The CCF controller acts a single point of integration with vCenter through a vCenter extension developed using vCenter APIs. It gets notified of events from the vCenter, upon which it performs the corresponding fabric operations and eliminates the need for any manual configuration of the physical network.

Automatic Host Detection and Link Aggregation

CCF can automatically detect where ESXi nodes connect to the fabric and create interface-groups for the node. This makes ESXi host addition/removal on vSphere cluster zero touch for the network.

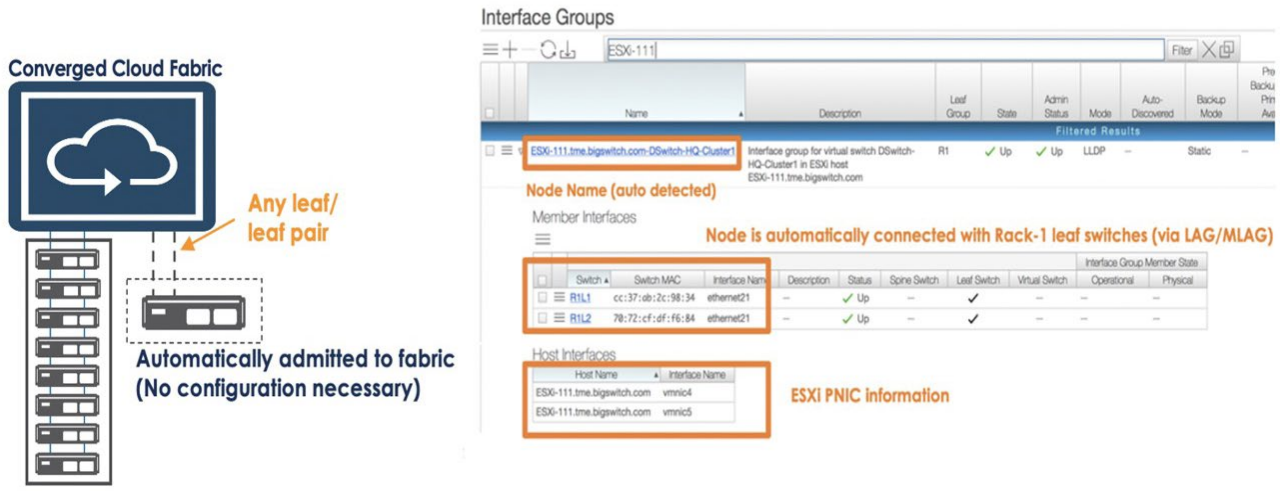

Figure 3: CCF Host Automation

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.

E-VPC Automation for L2 Network Creation & VM Learning

As part of the application deployment process, vCenter creates, modifies or deletes VMware virtual switch port-groups. CCF controller gets notified of these events and automatically creates, modifies or deletes the corresponding CCF Layer-2 network segments. A newly created VM in the vCenter is dynamically learned by CCF as an endpoint. The CCF controller then enables segment membership for the designated LAG connected to the corresponding host and programs the forwarding tables. If the added VM is the first one on the host for the vCenter port-group, the corresponding CCF segment is provisioned on the host interface.

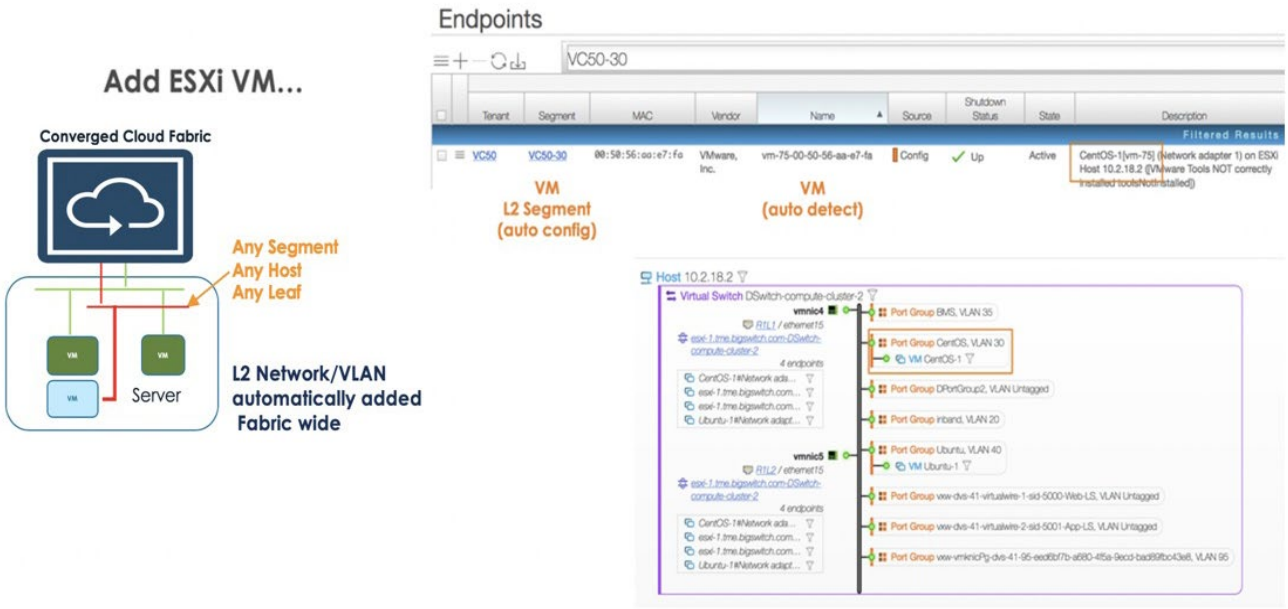

Figure 4: CCF L2 Network Automation

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.

Network Policy Migration for vMotion/DRS

After an application is deployed and VMs are up and running, vMotion is a powerful capability to seamlessly move VMs from one host to another. When vMotion is initiated in vCenter, The CCF controller gets notified of the new location of the VM information and migrates the network policies dynamically and updates the forwarding tables with the new information. If the VM is the first one in that VMware port-group on the new host, CCF will automatically provision the VLAN on the new host's LAG port.

Similarly, if the VM is the last one on the VMware port-group on the old host, CCF will prune the port group from the old host after the vMotion is complete

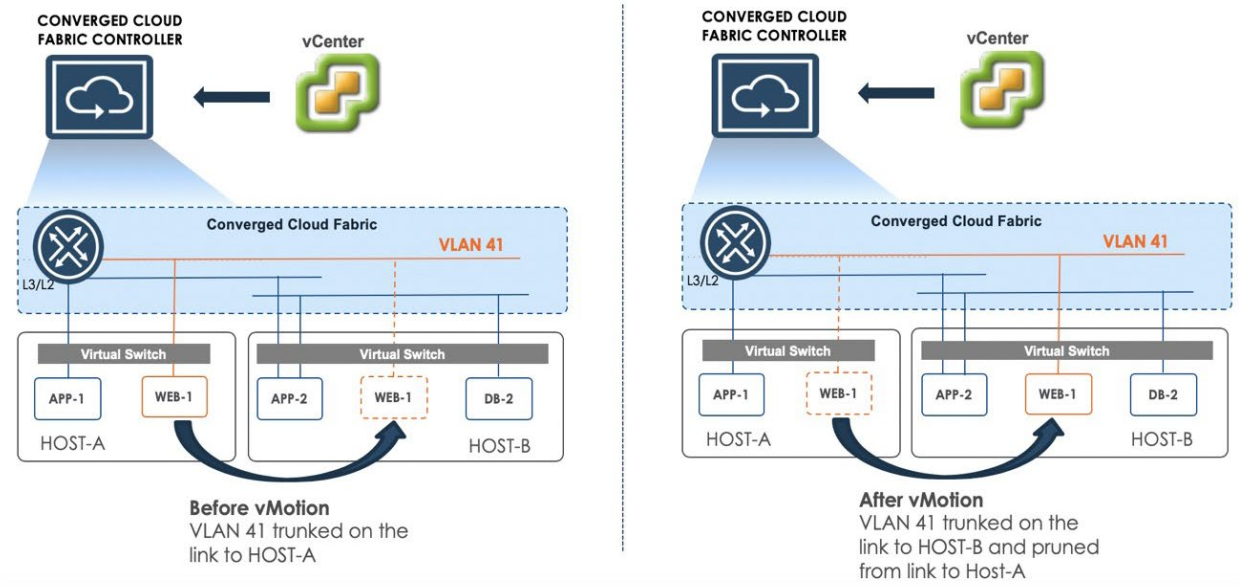

Figure 5: CCF Network Policy Migration

Contextual Visibility, Analytics and Troubleshooting with vSphere and CCF

The key challenge with networking in virtual environments is the lack of visibility on the end-to-end connectivity between VM endpoints across both the virtual and physical networks. The CCF/vSphere solution offers significant benefits to data center administrators by providing advanced end-to-end visibility and enhanced troubleshooting capabilities, all exposed through a single pane of the glass, the CCF GUI.

Contextual Visibility

The CCF GUI presents all the VM-related information that it learned from the vCenter. The display includes all of the hosts, their vmNICs, and the physical fabric interfaces to which the hosts are connected. It also includes the VM endpoint information including the name, logical segment, IP, MAC, and LAG. All of this information helps the admin get a quick understanding of fabric configuration to begin a troubleshooting session Provide VM, Host and VDS level visibility to the Network Admin from CCF GUI

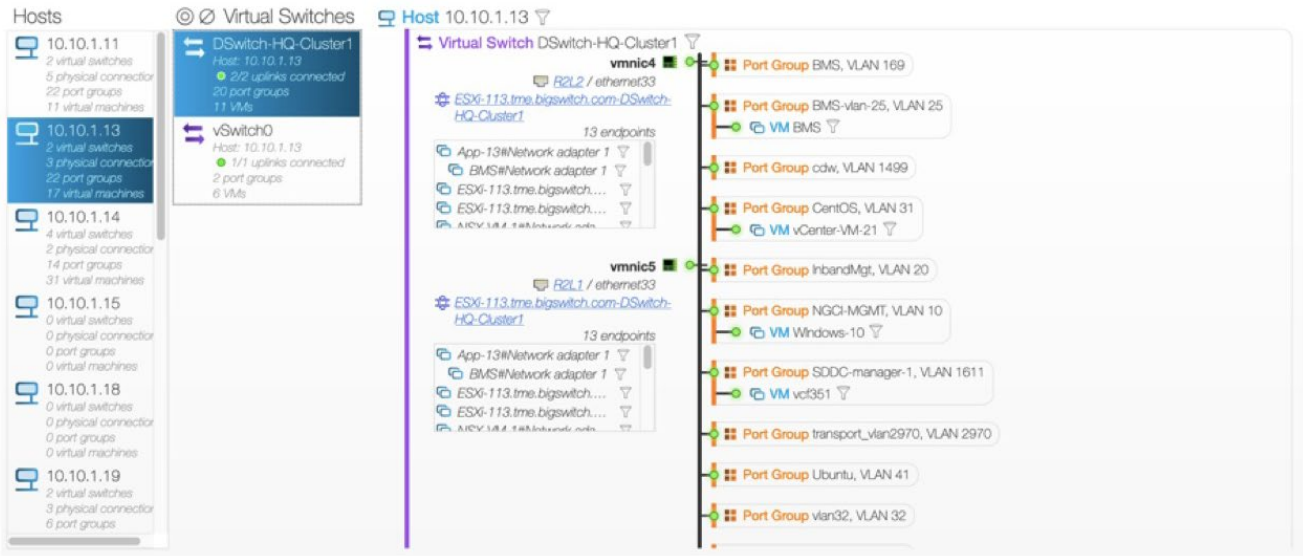

Figure 6: CCF Contextual Visibility

Fabric Analytics for VMware Networking

CCF provides advanced fabric analytics for VMware networking with a graphical representation of all VM-related information (name, creation time, pNIC info, port-groups), and time series of events related to VMs. This is a tremendous asset for troubleshooting as the administrator can get to the details of specific events for one or more VMs or can zoom into a timeslot to obtain all events that may have occurred during the period.

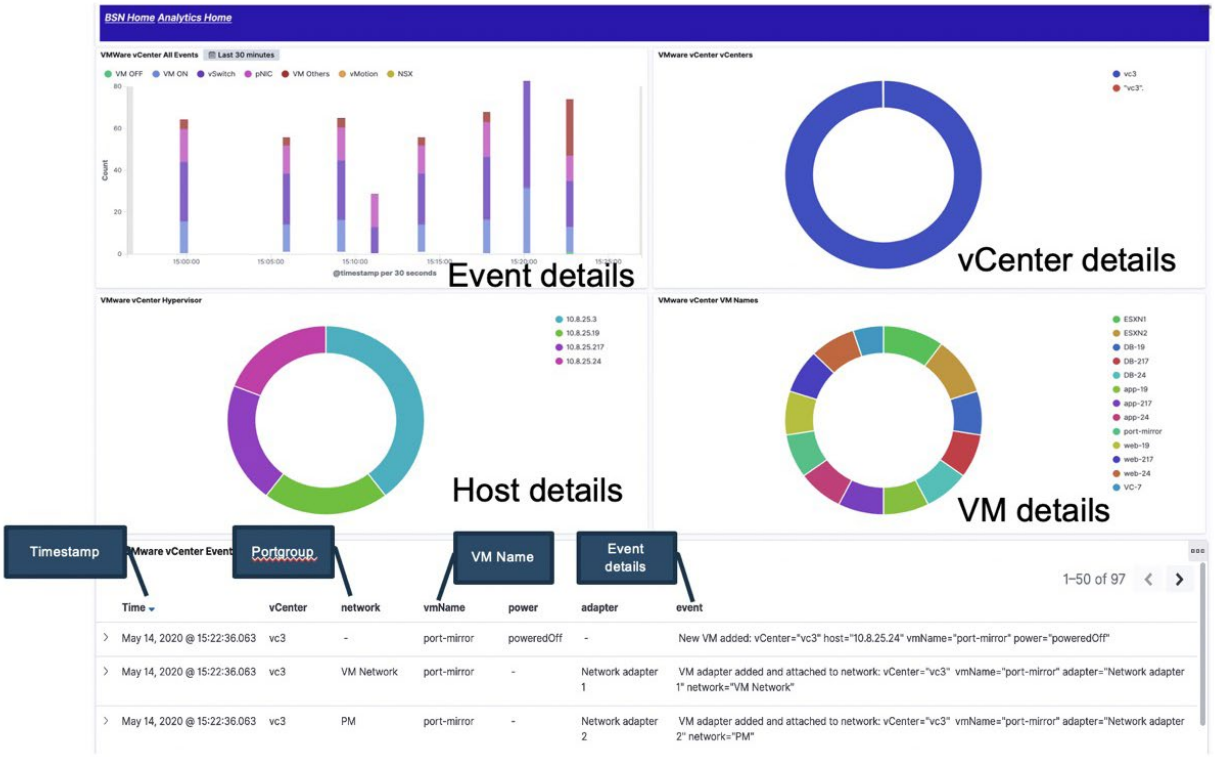

Figure 7: CCF Fabric Analytics

One Click Troubleshooting

VM-to-VM and vmkernel-to-vmkernel traffic visibility across the virtual and physical network can be simulated using CCF's test path feature which displays on the controller both physical and logical path taken by the traffic from one VM to another. This level of visibility to traffic, which cannot be achieved with box-by-box networking, helps rapidly determine if an application issue is network-related versus compute-related without going through tedious trouble ticket processes.

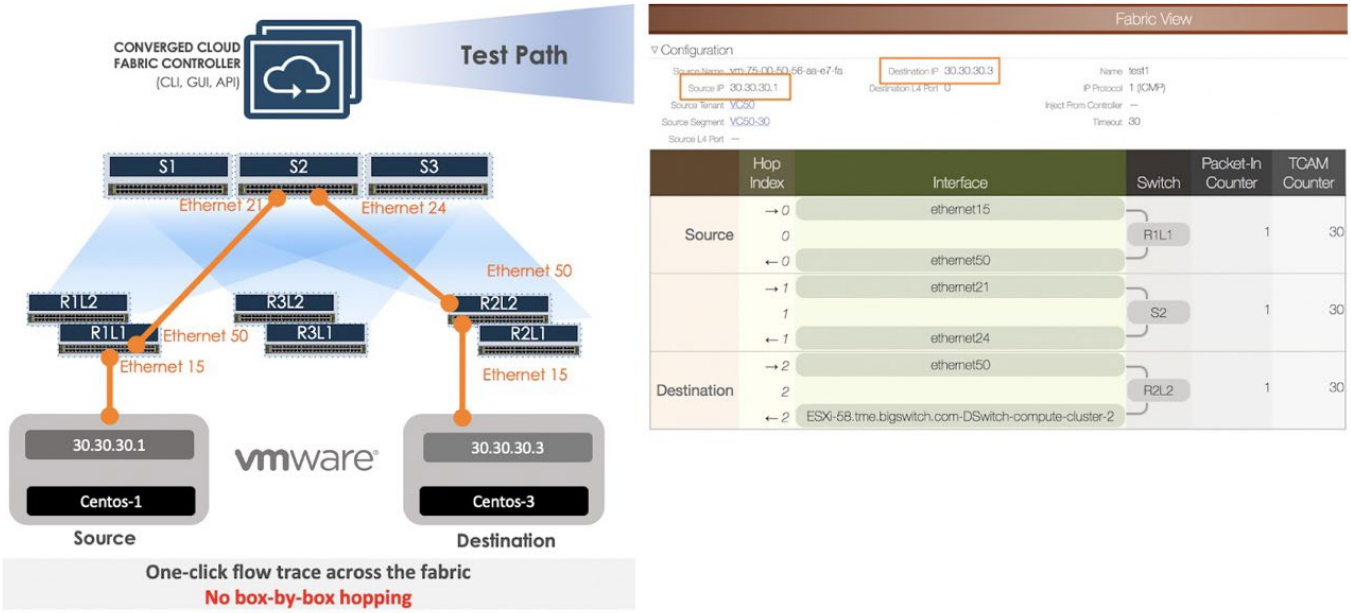

Figure 8: CCF One Click Troubleshooting

Summary

With Arista Converged Cloud Fabric solution for vSphere, Arista Networks offers the most comprehensive, flexible and highly automated solution in the industry.

		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