

# Converged Cloud Fabric™

## Self-Service Automated Fabric with Cloud-Style VPCs ON-PREM

### Converged Cloud Fabric Benefits

- Cloud on-prem experience with Enterprise VPC (E-VPC) for multi-tenant self-service networking
- Agility Boost for New Service Enablement with Network Automation for private cloud & HCI, including VMware (vSphere, NSX, vSAN, VCF), VXRail, Nutanix
- Faster Change Management with Zero-Touch Fabric Operations and SDN Controls
- Rapid Fabric-wide Upgrades with One-step Workflow
- One-click Troubleshooting with Fabric Trace, Policy (ACL) Trace and Contextual Analytics

Get hands-on experience with our offering — register for a free online trial at: [ccf-labs.arista.com](http://ccf-labs.arista.com)

Contact our sales team at: [ccf-sales@arista.com](mailto:ccf-sales@arista.com)

### Converged Cloud Fabric - Overview

Converged Cloud Fabric (CCF)™ is an automated fabric built with cloud networking design principles. CCF leverages public cloud-style VPC/VNet constructs on-prem to deliver a Network-as-a-Service operational model. CCF automates networking for multiple private cloud platforms, enabling the network to operate at the speed of VMs and Containers. With built-in analytics and telemetry, CCF provides real-time contextual visibility across the fabric and one-click troubleshooting workflows. With CCF, NetOps, DevOps and CloudOps teams can effectively collaborate, and rapidly on-board applications and tenants.

### SDN Software Meets VPCs ON-PREM

- To bring public cloud's self-service experience to on-prem data center networks, CCF has implemented AWS-style Virtual Private Cloud (VPC) in the form of Enterprise VPC (E-VPC). Similar to public cloud's VPC/VNet-based logical networking, CCF's E-VPCs decouple logical network policies from underlying network hardware attributes (such as port, switch, rack, VLAN, VRF). With E-VPCs, mainstream IT organizations can deliver cloud-style Network-as-a-Service and zero-touch operational experience for all private cloud platforms through built-in network automation including VMware vSphere/NSX/vSAN, DellEMC VxRail HCI, Nutanix HCI, Microsoft Hyper-V, Kubernetes containers and OpenStack. Additionally, E-VPC's built-in multi tenancy and delegated administration are ideal for DevOps/Cloud teams for programmatic automation. Mainstream IT organizations get cloud-style experience, infrastructure and economics, and are finally free of insurmountable complexity and vendor lock-in of traditional box-by-box networking.
- Software Defined Networking (SDN) fabric architecture refers to a separation of the network's data and control plane, followed by centralization of the control plane functionality. In practice, it implies that the network's policy plane, management plane, and much of the control plane are externalized from the hardware device itself using an SDN controller, with few on-device off-load functions for scale and resiliency. The network state is centralized but hierarchically implemented, instead of being fully distributed on a box-by-box basis across access and aggregation switches.
- Controller-based designs not only bring agility via centralized programmability and automation, but they also streamline fabric designs (e.g. leaf-spine L2/L3 Clos) that are otherwise cumbersome to implement and fragile to operate in a box-by-box design.

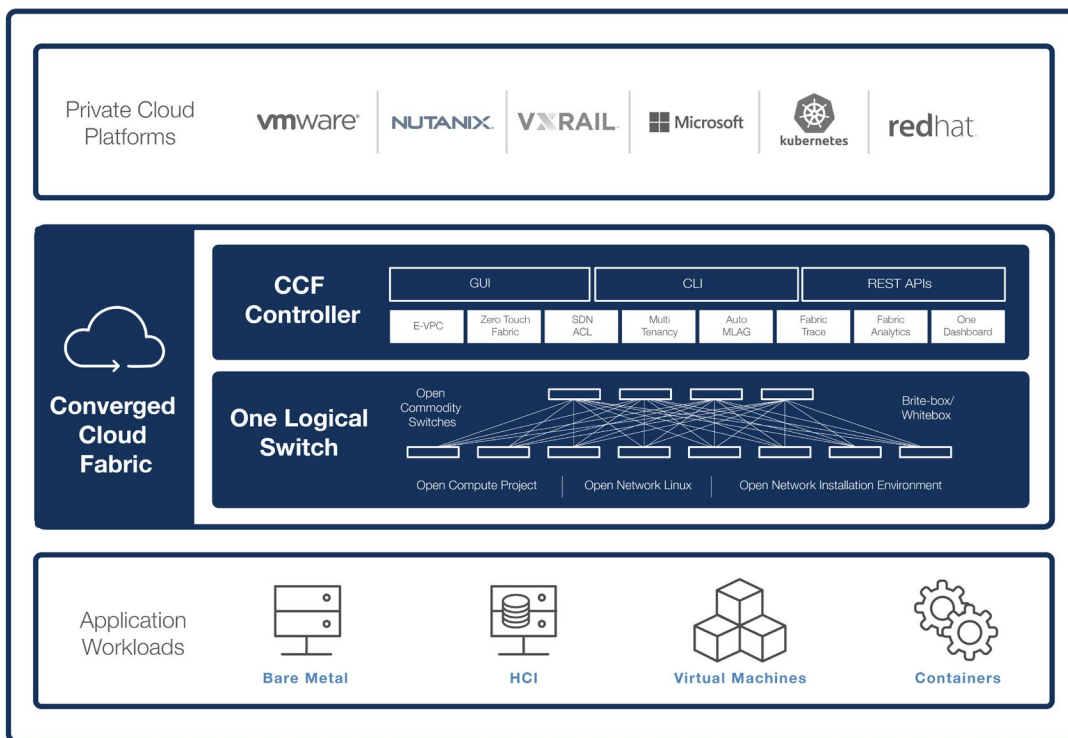


Figure 1: Converged Cloud Fabric (One Logical Switch)

The CCF architecture consists of a physical switching fabric, which is based on a leaf-spine Clos architecture. Leaf and spine switches running Switch Light™ Operating System form the individual nodes of this physical fabric. Intelligence in the fabric is hierarchically placed—most of it in the CCF Controller (where configuration, automation and troubleshooting occur), and some of it off-loaded to Switch Light for resiliency and scale-out.

### Converged Cloud Fabric System Components

- **Converged Cloud Fabric Controller** — a centralized and hierarchically implemented SDN controller available as an HA pair of hardware appliances for high availability.
- **Open Networking Leaf and Spine Switch Hardware** — the term ‘open networking’ (whitebox or britebox) refers to the fact that the Ethernet switches are shipped without an embedded networking OS. The merchant silicon networking ASICs used in these switches are the same as used by most incumbent switch vendors and have been widely deployed in production in hyperscale data center networks. These bare metal switches ship with Open Network Install Environment (ONIE) for automatic and vendor-agnostic installation of third-party network OS. A variety of switch HW configurations and vendors are available on the CCF hardware compatibility list.
- **Switch Light Operating System** — a light-weight open networking switch OS purpose built for SDN

- **Converged Cloud Fabric Automation Plugins (optional)**

— Converged Cloud Fabric has built-in E-VPC network automation for private cloud platforms including VMware vSphere/NSX/vSAN, DellEMC VxRail HCI, Nutanix HCI, Microsoft Hyper-V, Kubernetes containers and OpenStack.

### Deployment Solutions

CCF is designed from the ground up to satisfy the requirements of physical, virtual or combination of physical and virtual workloads. It supports a wide variety of data center and private cloud use cases, including:

- VMware SDDC workloads (vSphere, NSX, Virtual SAN and VIO)
- Converged and Hyper-converged Systems (HCI) with Nutanix and Dell EMC VxRail
- Microsoft Hyper-V workloads
- OpenStack including NFV
- Kubernetes orchestrated containerized workloads
- Virtual desktop infrastructure (VDI) workloads

Example Scenario	Supported Workloads	Leaf Switch Configuration	Spine Switch Configuration
<b>Private Cloud—Typical data center pod deployments</b>	1G, 10G, 25G	48X10G + 6x40G, 32x100G, 48x25G + 6x100G, 48x25G + 8x100G	32x40G, 32x100G, 64x100G
<b>Cost Optimized Fabric (leverage existing cable infrastructure)</b>	1G, 10G	48X10G + 6x40G	48X10G + 6x40G
<b>Software Defined Storage / Converged Data Analytics</b>	25G, 40G	32x40G, 32x100G	32x40G, 32x100G, 64x100G
<b>Dense 10G, 25G Compute (using splitter cables)</b>	10G, 25G	32x40G, 32x100G	32x40G, 32x100G, 64x100G
<b>High performance 40G storage array and 10G workloads (using splitter cables)</b>	10G, 25G, 40G	48X10G + 6x40G, 32x40G, 32x100G	32x40G, 32x100G, 64x100G

Figure 2: Example CCF Deployment Scenarios

### Converged Cloud Fabric Benefits

#### Centralized Controller Reduces Management Consoles By Over 60:1

With configuration, automation and most troubleshooting done via the CCF Controller, the number of management consoles involved in provisioning new physical capacity or new logical apps goes down dramatically. For example, in a 32 rack pod with dual leaf switches and four spine switches, a traditional box-by-box network design would have 68 switch management consoles. The Converged Cloud Fabric design has only one—the controller console—that performs the same functions. The result is massive time savings, reduced error rates and simpler automation designs. As a powerful management tool, the controller console exposes a web-based GUI, a traditional networking-style CLI and REST APIs.

#### Streamlined Configuration, Enabling Rapid Innovation

In the CCF design, configuration in the CLI, GUI or REST API is based on the concept of logical tenants. Each tenant has administrative control over a logical L2/L3/policy design that connects the edge ports under the tenant’s control. The Converged Cloud Fabric controller has the intelligence to translate the logical design into optimized entries in the forwarding tables of the spine, leaf and vleaf.

#### Open Networking Switch Hardware Reduces CapEx Costs By Over 50%

By adding up hardware, software, maintenance and optics/cables, a complete picture of the hard costs over three years shows that the savings are dramatic.

#### Built-in Orchestration Support Streamlines DC Operations

CCF Controller natively supports integration with popular private cloud platforms—VMware vSphere/NSX/vSAN, DellEMC VxRail HCI, Nutanix HCI, Microsoft Hyper-V, Kubernetes containers and OpenStack—through a single programmatic interface. This is tremendously simpler and scalable compared to traditional box-by-box complexity networking which demands an exponentially larger number of programmatic interactions with CMPs. Data center admins benefit from streamlined application deployment workflows, enhanced analytics and simplified troubleshooting across physical and virtual environments.




Figure 3: E-VPC Automation for Cloud Management Platforms



Figure 4: CCF Graphical User Interface (GUI)

### SDN Fabric Enables Deep Visibility and Resilience for OpenStack Networks

The CCF OpenStack Neutron plugin for L2 networking provides resiliency necessary for production-grade OpenStack deployments. The CCF Controller acts as the single pane for provisioning, troubleshooting, visibility and analytics of the entire physical network environment. This enables data center operators to deploy applications rapidly, simplifies operational workflows and provides immediate root-cause analysis when application performance issues arise.

### Scale-out (Elastic) Fabric

CCF's flexible, scale-out design allows users to start at the size and scale that satisfies their immediate needs while future proofing their growth needs. By providing a choice of hardware and software solutions across the layers of the networking stack and pay-as-you-grow economics, starting small scale and growing the fabric gradually instead of locking into a fully integrated proprietary solution, provides a path to a modern data center network. Once new switches are added, the controller adds those switches to the fabric and extends the current configuration hence reducing any error that may happen otherwise. Customers take advantage of one-time configuration of the fabric.

### DC-grade Resilience

CCF provides DC grade resiliency that allows the fabric to operate in the face of link or node failures as well as in the rare situation when the controller pair is unavailable (headless mode). Swapping a switch (in case of HW failure or switch repurpose) is similar to changing a line card in a modular chassis. After re-cabling and power up, the switch boots up by downloading the right image, configuration and forwarding tables. Additionally, the CCF Controller coordinates and orchestrates the entire fabric upgrade ensuring minimum fabric down time. These functionalities further enhance fabric resiliency and simplify operations.

### Using CCF: A 3-Tier Application Example

CCF supports a multi-tenant model, which is easily customizable for the specific requirements of different organizations and applications. This model increases the speed of application provisioning, simplifies configuration, and helps with analytics and troubleshooting. Some of the important terminology used to describe the functionality include:

- E-VPC Tenant — A logical grouping of L2 and/or L3 networks and services.
- Logical Segment — An L2 network consisting of logical ports and end-points. This defines the default broadcast domain boundary.
- Logical Router — A tenant router providing routing and policy enforcement services for inter-segment, inter-tenant, and external networks.
- External Core Router — A physical router that provides connectivity between Pods within a data center and to the Internet.
- Tenant Services — Services available to tenants and deployed as dedicated or shared services (individually or as part of a service chain).

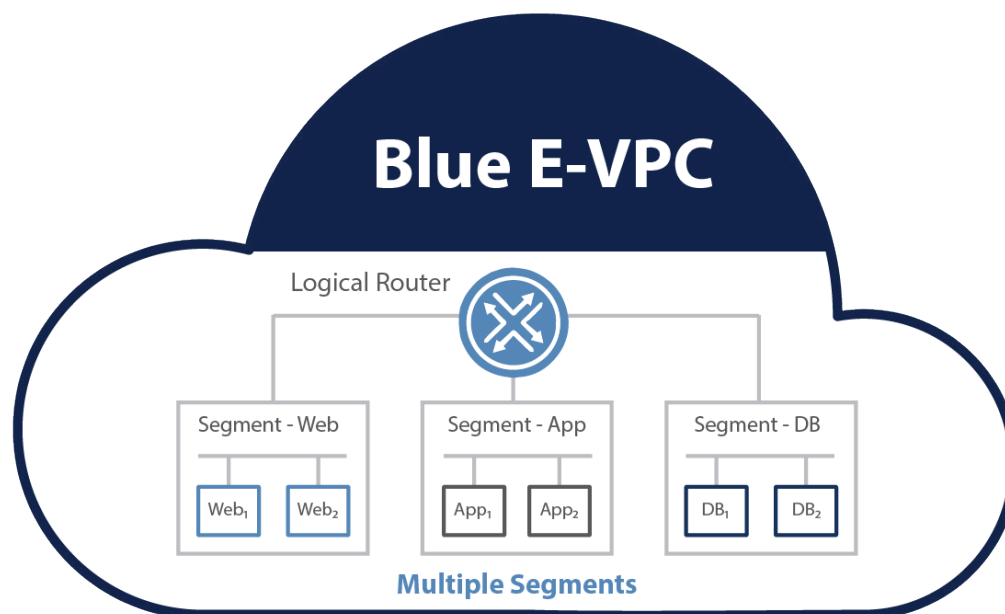


Figure 5: CCF Logical Topology

### E-VPC (Tenant)

In the most common scenario, end consumers or tenants of the data center infrastructure deal with a logical network topology that defines the connectivity and policy requirements of applications. As an illustrative example, the canonical 3-tier application in Figure 5 shows various workload nodes of a tenant named “BLUE.” Typically, a tenant provisions these workloads using orchestration software such as OpenStack, VMware vSphere, or CCF Controller GUI/CLI, directly. As part of that provisioning workflow, the CCF Controller seamlessly enables the logical topology onto the physical and virtual switches.

### Mapping Logical to Physical

The BLUE Tenant has three logical network segments. Each of the three segments represents the broadcast domain for the 3-tiers—Web, App, and Database. Let’s say in this example, Web1,2 and App1,2 are virtualized workloads but DB1,2 is comprised of physical workloads. Following the rules defined by the data center administrator, the orchestration system provisions requested workloads across different physical nodes within the data center. As an example, the logical topology shown in Figure 5 could be mapped on the pod network. The CCF Controller handles the task of providing optimal connectivity, between these loads dispersed across the pod, while ensuring tenant separation and security.

In order to simplify the example, we only show racks that host virtualized and physical workloads in the figure above, but similar concepts apply for implementing tenant connectivity to external router and chaining shared services.

An illustrative sample set of entries in various forwarding tables highlight some of the salient features of CCF described in earlier sections.

- L3 routing decision is made at the first hop leaf switch (distributed virtual routing—no hair-pinning)
- L2 forwarding across the pod without special fabric encapsulation (no tunneling)
- Full load-balancing across the various LAG links (leaf and spines)
- Leaf/Spine mesh connectivity within the physical fabric for resilience

```
! evpc-tenant
evpc-tenant BLUE
logical-router
route 0.0.0.0/0 evpc-tenant system
interface segment web
ip address 10.1.1.254/24
interface segment db
ip address 10.1.2.254/24

segment web
member interface-group IG-101 vlan 101
segment db
member interface-group IG-102 vlan 102

evpc-tenant DMZ
logical-router
ospf
ospf-interface segment DMZ
area 0
interface segment DMZ
ip address 20.1.1.254/24
```

Figure 6: Application Centric Configuration

Feature	Description / Benefit
<b>Enterprise VPC (E-VPC)</b>	<p>Converged Cloud Fabric provides VPC-style logical networking called Enterprise VPC, delivering on-prem VPC. Cloud-style E-VPC enables multi-tenant network-as-a-service.</p> <ul style="list-style-type: none"> <li>• E-VPC makes networking as simple as VPC/VNETs on public clouds</li> <li>• Simplifies L2/L3 provisioning and hardware constructs (ports, VLANs, VRFs)</li> <li>• Delegated administration with RBAC control</li> </ul>
<b>Zero Touch Fabric (ZTF)</b>	<p>ZTF enables complete control and management of physical switches within CCF without manually interacting with the switches. It tremendously simplifies day-to-day network operations:</p> <ul style="list-style-type: none"> <li>• Auto-configuration and auto-upgrade of Switch Light OS</li> <li>• Automatic topology updates and event notifications based on fabric link state changes</li> <li>• Auto-scaling of the fabric—adding or removing nodes and/or links within the fabric requires no additional configuration changes on the controller</li> </ul>
<b>Fabric LAG</b>	<p>Fabric LAG combines the underlying LAG functionality in switching ASICs with the centralized visibility of the SDN controller to create a highly resilient and efficiently balanced fabric. Compared to spanning tree protocols or even traditional MLAG/ECMP based approaches to multi-path fabric formation, Fabric LAG technology enables significantly reduced convergence time on topology changes and dramatically reduced configuration complexity.</p>
<b>Fabric Sync (Controller RIB - FIB)</b>	<p>Fabric Sync intelligently synchronizes Controller Routing Information Base (RIB) with fabric node's Forwarding Information Base (FIB). During a topology change, only delta updates are synced across impacted switches. Fabric Sync ensures strong RIB-FIB consistency, as it is the single point of control for maintaining all forwarding and associated policy tables.</p>
<b>Resilient Headless Mode</b>	<p>In situations when both controllers are unreachable, fabric nodes are considered to be running in Headless mode. In this mode, all provisioned services continue to function as programmed prior to entering the Headless mode. Additionally, multiple levels of redundancy enable a highly resilient and self-healing fabric even during headless mode.</p>
<b>SDN-Managed Fabric (GUI, CLI &amp; REST APIs)</b>	<p>Converged Cloud Fabric Controller provides single pane of glass for entire fabric.</p> <ul style="list-style-type: none"> <li>• Administrators can configure, manage, debug or troubleshoot, and upgrade the fabric nodes using CLI, GUI, or REST API.</li> <li>• REST APIs, CLI and GUI have application and tenant awareness.</li> </ul> <p>Single Pane of Glass fabric management enhances operational simplicity by providing a centralized dashboard for fabric management as well as quick and easy access to troubleshooting, analytics and telemetry information. Additionally, it provides simplified workflow for network operators and administrators.</p>



Feature	Description / Benefit
<b>Fabric Analytics</b>	Fabric Analytics provides Advanced Multi-node Troubleshooting, Analytics & Telemetry in the Converged Cloud Fabric solution. Data sets in Fabric Analytics can also be streamed to remote data collection systems like Splunk/Kafka, etc.
<b>API-first Fabric</b>	Converged Cloud Fabric Controller is highly programmable due to its “API-first” design principle and can be implemented as a closed loop feedback system. For example, security applications can dynamically detect threats and program the CCF controller for mitigation. The CCF GUI and CLI utilize the underlying REST APIs—hence are by definition consistent and hardened.
<b>Tenant-aware Fabric</b>	Converged Cloud Fabric provides built-in multi-tenancy via tenant-aware configurations, tenant separation, and fine-grain inter-tenant access control. Configuration in the CLI, GUI or REST API is based on the concept of logical tenants.
<b>Service-aware Fabric</b>	Converged Cloud Fabric supports L3 virtual and physical service insertion and service chaining. Services can be shared across tenants or dedicated to a specific tenant.
<b>L2 Features</b>	<ul style="list-style-type: none"> <li>• Layer 2 switch ports and VLAN trunks</li> <li>• IEEE 802.1Q VLAN encapsulation</li> <li>• Support for up to 4K VLANs (i.e. 4K Logical Segments)</li> <li>• MAC address based segmentation</li> <li>• BPDU Guard</li> <li>• Storm Control</li> <li>• MLAG (up to 16 ports per LAG)</li> <li>• 3,800 IGMP Groups</li> <li>• IGMP Snooping</li> <li>• Static Multicast Group</li> <li>• Link Layer Discovery Protocol (LLDP)</li> <li>• Link Aggregation Control Protocol (LACP): IEEE 802.1AX</li> <li>• LACP Fallback Mode (Dynamic membership management for server PXE booting)</li> <li>• Jumbo frames on all ports (up to 9216 bytes)</li> <li>• VLAN Translation</li> <li>• Primary / Backup Interface</li> <li>• VXLAN Support</li> <li>• Preserve VLAN (Q-in-Q)</li> </ul>

Feature	Description / Benefit
<b>L3 Features</b>	<ul style="list-style-type: none"> <li>• Layer 3 interfaces: Routed ports, Switch Virtual Interface (SVI), Distributed Gateway</li> <li>• Multiple IP-Subnet Support per Segment/SVI</li> <li>• Support for up to 90K IPv4 host prefix, 14K IPv6 host prefix (i.e. Endpoints)</li> <li>• Support for 1K Virtual Routing and Forwarding (VRF) entries (i.e. 1K Logical Routers)</li> <li>• 1K Tenants, 1K E-VPCs</li> <li>• Static Route, BGP (IPv4, IPv6), OSPF (IPv4), OSPFv3 (IPv6)</li> <li>• Static and dynamic route exchange between E-VPCs</li> <li>• 68K IPv4 routes, 8K IPv6 routes</li> <li>• Up to 16 ways Equal-Cost Multipathing (ECMP)</li> <li>• 1K Equal-Cost Multipathing (ECMP) groups</li> <li>• 6K flexible ACL entries</li> <li>• Policy-Based Routing</li> <li>• Multicast Routing, PIM-SM</li> <li>• ACL: Routed ACL with Layer 3 and 4 options to match ingress ACL</li> <li>• Jumbo frame support (up to 9216 bytes)</li> <li>• DHCP/DHCPv6 relay</li> <li>• IP Helper</li> <li>• NAT/PAT support</li> <li>• BFD support for BGP</li> </ul>
<b>QoS</b>	<ul style="list-style-type: none"> <li>• Layer 2 IEEE 802.1p (class of service [CoS])</li> <li>• Source segment or IP DSCP based Classification</li> <li>• Tenant/Segment based classification</li> <li>• DWRR based egress queuing</li> <li>• CoS based marking</li> <li>• PFC and DCBX</li> <li>• IP address/subnet based QoS classification</li> </ul>
<b>High Availability</b>	<ul style="list-style-type: none"> <li>• Controller HA</li> <li>• Headless mode (fabric forwards traffic in absence of Controller)</li> <li>• Redundant Spine</li> <li>• Redundant Leaf</li> <li>• Redundant Links</li> <li>• Controller cluster with single Virtual IP</li> <li>• Support for redundant out-of-band management switch</li> </ul>



Feature	Description / Benefit
<b>Security</b>	<ul style="list-style-type: none"> <li>• Ingress ACLs</li> <li>• Layer 3 and 4 ACLs: IPv4, Internet Control Message Protocol (ICMP), Transmission Control Protocol (TCP), User Datagram Protocol (UDP), etc.</li> <li>• ACLs on controller management interface</li> <li>• ACL logging (IPv4 only)</li> <li>• Time based lockout of user accounts</li> <li>• Control Plane Policing (CoPP) or Rate Limiting</li> <li>• Custom TLS keys and Certs support for GUI</li> <li>• EAP-TTLS Support for Radius Authentication</li> <li>• Two Factor Authentication</li> <li>• Restrict Cipher Suites</li> <li>• FIPS 140-2 compatibility</li> <li>• PKI Certificate Revocation</li> <li>• TLS v1.3 support</li> </ul>
<b>OpenStack Network</b>	<p>Provides Fabric Automation and Visibility for Openstack, including:</p> <ul style="list-style-type: none"> <li>• E-VPC Automation for ML2 Driver Mechanism</li> <li>• Dynamic Provisioning of the CCF Fabric</li> <li>• LBaaS support (network automation driven through OpenStack)</li> <li>• SR-IOV Integration</li> <li>• OVS DPDK Integration</li> <li>• OpenStack Cluster Visibility</li> </ul>
<b>VMware vSphere Network Automation</b>	<p>Provides Fabric Automation and Visibility including:</p> <ul style="list-style-type: none"> <li>• Auto Host Detection &amp; LAG Formation</li> <li>• E-VPC Automation for L2 Network Creation &amp; VM Learning</li> <li>• Network Policy Migration for vMotion/DRS</li> <li>• VM-level Visibility (VMname, vMotion)</li> <li>• VM-to-VM Troubleshooting (Logical &amp; Physical)</li> <li>• L3 configuration via vSphere web-client plugin</li> <li>• Test Path visibility through vCenter</li> <li>• Multiple E-VPC tenants per vCenter</li> </ul>

Feature	Description / Benefit
<b>VMware NSX Underlay Network Automation</b>	Close the overlay/underlay gap for visibility and troubleshooting. Features include: <ul style="list-style-type: none"> <li>• Auto Host Detection &amp; LAG Formation</li> <li>• E-VPC Automation for physical network creation</li> <li>• Underlay Troubleshooting – VM-to-VM &amp; TEP-to-TEP connectivity</li> <li>• Underlay Visibility through Fabric Analytics (VM-name, VXLAN ID, Logical Switch)</li> <li>• NSX Hardware VTEP support</li> <li>• NSX-T (Geneve &amp; VLAN)</li> </ul>
<b>VMware vSAN Network Automation</b>	Provides Fabric Automation and Visibility for vSAN, including: <ul style="list-style-type: none"> <li>• Auto-detection and LAG formation for vSAN node</li> <li>• E-VPC automation for creation of vSAN transport network</li> <li>• vSAN cluster communication troubleshooting for unicast and multicast</li> <li>• Simplified Layer 2 / Layer 3 multicast deployment for vSAN transport</li> <li>• vSAN Analytics</li> </ul>
<b>Nutanix Network Automation</b>	Provides Fabric Automation and Visibility for Nutanix, including: <ul style="list-style-type: none"> <li>• Automatic Host Bootstrapping</li> <li>• Auto Host Detection &amp; LAG Formation (support all AHV load-balancing modes)</li> <li>• E-VPC automation for physical L2 Network Creation &amp; VM Learning</li> <li>• Auto L3 Network Creation &amp; Distributed Logical Routing</li> <li>• Network Policy Migration for VM Migrations</li> <li>• AHV Networking and VM-level Visibility (VMname, Host Information, Physical Connectivity)</li> <li>• VM-to-VM Troubleshooting with Test-Path (Logical &amp; Physical)</li> <li>• Multiple tenants per AHV Cluster</li> <li>• Nutanix Flow Microsegmentation visibility</li> </ul>
<b>Hyper-V Network Automation</b>	Provides Fabric Automation and Visibility for Hyper-V, including: <ul style="list-style-type: none"> <li>• Auto-host detection via Microsoft LLDP agent on nodes</li> <li>• E-VPC automation for physical network configuration</li> <li>• Network policy migration</li> <li>• VM-level visibility (VM MAC &amp; VM IP address)</li> <li>• VM-to-VM Troubleshooting (Logical &amp; Physical)</li> </ul>

Feature	Description / Benefit
<b>Container Network Automation</b>	Provides Fabric Automation and Visibility for Kubernetes, including: <ul style="list-style-type: none"> <li>• Auto-Host Detection &amp; MLAG Formation</li> <li>• E-VPC automation of physical L2 network creation for Container Network through Kubernetes Custom Resource Definitions (CRDs)</li> <li>• Multiple K8s cluster support and K8s micro-services visibility</li> <li>• Multiple Kubernetes CNI support including Calico, Flannel, Cilium &amp; Canal</li> <li>• Container-level visibility (Container name, Namespace, Node Name, Status, IP)</li> <li>• Micro-services visibility including Microservices to Endpoints/IP &amp; Port mappings</li> <li>• K8S Node-to-Node troubleshooting (Logical &amp; Physical)</li> </ul>
<b>Multi-Orchestration Support</b>	Support Multiple Enterprise Virtual Private Cloud (E-VPCs) on single CCF Fabric
<b>Inter-Pod Connectivity</b>	<ul style="list-style-type: none"> <li>• L3 — Using Static Route, OSPF &amp; BGP</li> <li>• L2 — Dark Fiber</li> <li>• L2 — VXLAN</li> </ul>
<b>MIBs</b>	Documented in a separate MIBs document
<b>Ethernet Standards</b>	<ul style="list-style-type: none"> <li>• IEEE 802.1p: CoS prioritization</li> <li>• IEEE 802.1Q: VLAN tagging</li> <li>• IEEE 802.3: Ethernet</li> <li>• IEEE 802.3ab Gigabit Ethernet (1000BASE-T) or breakout</li> <li>• IEEE 802.3ae: 10 Gigabit Ethernet</li> <li>• IEEE 802.3by: 25 Gigabit Ethernet</li> <li>• IEEE 802.3ba: 40 Gigabit Ethernet and 100 Gigabit Ethernet</li> <li>• IEEE 802.1AB LLDP</li> <li>• IEEE 802.3ad Link Aggregation (LACP)</li> <li>• IEEE 802.1Qbb Priority-Based Flow Control (PFC)</li> <li>• IETF: IPv4 &amp; IPv6</li> </ul>
<b>Analytics Node Integration</b>	Enables network performance monitoring and simplifies app vs network troubleshooting Features include: <ul style="list-style-type: none"> <li>• IPv4 End-point discovery (using ARP &amp; DNS)</li> <li>• IPv6 End-point discovery (using ICMPv6 neighbor discovery message)</li> <li>• OS Fingerprinting (using DHCP)</li> <li>• End-to-end flow analysis (using sflow)</li> </ul>

Feature	Description / Benefit
<b>Switch Platform Support</b>	<p><b>Arista Switch Platforms</b> Support Arista 10G/25G and 40G/100G platforms:</p> <ul style="list-style-type: none"> <li>• 48x25G + 12x100G</li> <li>• 48x25G + 8x100G</li> <li>• 48x10GT + 8x100G</li> <li>• 32x100G</li> <li>• 64x100G</li> </ul> <p><b>3rd-Party Switch Platform Support</b> Support for DellEMC 10G, 25G, 40G and 100G platforms. The common supported switch configurations are:</p> <ul style="list-style-type: none"> <li>• 48x10G + 6x40G/4x100G</li> <li>• 48x10GT + 6x40G/4x100G</li> <li>• 48x25G + 6x100G/8x100G</li> <li>• 32x40G</li> <li>• 32x100G</li> <li>• 64x100G</li> </ul> <p>For the complete list of supported switch vendors/configurations and optics/cables included in the Converged Cloud Fabric Hardware Compatibility List (HCL), please contact the Arista Sales Team (sales@arista.com).</p>
<b>Fabric Management</b>	<ul style="list-style-type: none"> <li>• GUI (IPv4 / IPv6)</li> <li>• CLI (IPv4 / IPv6) — based console to provide detailed out-of-band management</li> <li>• Switch management using 10/100/1000-Mbps management through controller</li> <li>• Beacon LED (based on underlying switch)</li> <li>• Configuration synchronization</li> <li>• Configuration save and restore</li> <li>• Secure Shell Version 2 (SSHv2) — IPv4 / IPv6</li> <li>• Username and passwords authentication</li> <li>• TACACS+ / RADIUS — IPv4 / IPv6</li> <li>• Control Plane Security (CPSec) — Encrypted communication between Controllers and Physical / Virtual Switches</li> <li>• Syslog (4 servers) — IPv4 / IPv6, Syslog over TLS</li> <li>• SNMP v1, v2c and v3 — IPv4 / IPv6</li> <li>• sFlow support</li> <li>• SPAN with Policy/ACL</li> <li>• Fabric SPAN with Policy/ACL</li> <li>• Connected device visibility</li> <li>• Ingress and egress packet counters per interface, per segment, and per tenant</li> <li>• Network Time Protocol (NTP) — IPv4 / IPv6</li> <li>• Test Path — Enhanced Troubleshooting &amp; Visibility with logical and physical fabric views (VM &lt;--&gt; vLeaf &lt;--&gt; Leaf &lt;--&gt; Spine &lt;--&gt; Leaf &lt;--&gt; vLeaf &lt;--&gt; VM)</li> <li>• Fabric Analytics including telemetry and enhanced analysis</li> <li>• Prometheus support</li> </ul>
<b>Automation</b>	<ul style="list-style-type: none"> <li>• REST API (IPv4 / IPv6)</li> <li>• Ansible</li> <li>• CCF Terraform Provider</li> </ul>

### CCF Controller Appliance Specification

The CCF Controller can be deployed either as a physical appliance (production or lab deployment) or as a virtual machine appliance (for limited scale production or lab deployment).

#### CCF Controller — Physical Appliance Specification:

The CCF controller is available as an enterprise-class, 2-sockets, 1U rack-mount physical appliance designed to deliver the right combination of performance, redundancy and value in a dense chassis.



*Arista CCF Controller: DCA-CC-CDL*

Feature	Technical Specification*
<b>Controller</b>	<b>DCA-CC-CDL</b>
<b>Recommended for</b>	CCF P-fabric or P+V fabric
<b>Processor</b>	Intel Xeon 2 sockets (10 cores)
<b>Form Factor</b>	1U Rack Server
<b>Weight</b>	38.6 lbs
<b>Memory</b>	4 x 16GB
<b>Hard Drive</b>	2 x 1TB SATA (w/RAID support)
<b>Networking</b>	2 x 1Gb, 2 x 10Gb, 2 x 10Gbase-T
<b>Power</b>	Input Power: 302 watts Max Power : 403.1 watts Input Current: 2.7 Amps
<b>MTBF</b>	104,000 hours

\* Detailed environment information provided in CCF Hardware Guide.

### VM Appliance Specification

The Converged Cloud Fabric Controller VM Small is available as a Virtual Machine appliance for P or P+V fabric (for limited scale production or lab deployment).

Environment	CCF Controller VM Small	CCF Controller VM Medium
<b>VMware ESXi</b>	Version 6.5, 6.7, 7.0	Version 6.5, 6.7, 7.0
<b>Red Hat RHEL</b>	RHEL 7.2, 7.4, 7.5	RHEL 7.2, 7.4, 7.5
<b>vCPU</b>	6 vCPU	12 vCPU
<b>vMemory</b>	36 GB of Virtual Memory	46 GB of Virtual Memory
<b>HDD</b>	400GB HDD	400GB HDD
<b>vNIC</b>	4 vNICs	4 vNICs

Note: A VM's performance depends on many other factors in the hypervisor setup, and as such, we recommend using a hardware appliance for production deployments greater than three racks.

### Supported Workloads & Orchestration Systems

Feature	Technical Specification
<b>Physical Workloads</b>	Bare-metal server workloads
<b>Virtual Workloads</b>	VMware Integration with vSphere 6.5,6.7, 7.0, 8.0 Nutanix Integration with AOS 5.20x - 6.1..x Hyper-V Integration with Windows Server 2016, 2019, VMM 2016, VMM 2019, VMM1807 VMware Horizon View VDI Support any VM workload on CCF P Fabric even without orchestration integration (e.g. Xen, Hadoop) For OpenStack and Container integration, please refer tables below
<b>Cloud Orchestration</b>	OpenStack (Neutron ML2 driver, Neutron L3 Plugin) VMware VIO Container Integration—Container Network Interface (CNI) plugin support

## Openstack Integration

Hypervisor	Openstack
KVM	RHEL 8.1 & above (RHOSP 16)

## Container Integration

Distro	Container Orchestration	Bare Metal OS
Kubernetes Community	Kubernetes 1.14 - 1.16	CentOS 7.5 / Ubuntu 16.04

**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](mailto: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 , 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 © 2023 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. July 6, 2023