Hewlett Packard Enterprise

big switch
networks

Big Cloud Fabric

Next generation data center switching

At Hewlett Packard Enterprise, we help customers use technology to slash the time it takes to turn ideas into value. In turn, we transform industries, markets and lives. Open networking solutions from HPE free you from vendor lock-in. We also give you the ability to scale your cloud data center network to meet your business requirements, while using the resources that best suit your needs and lowering your costs. Big Cloud Fabric together with HPE Altoline strengthen HPE's growing commitment to open network.

Big Cloud Fabric™ is the next-generation data center switching fabric delivering operational velocity, network automation and visibility for software-defined data centers and cloud native applications, while staying within flat IT budgets. Using hyperscale-inspired networking principles, software controls and leaf/spine CLOS fabric delivered on open networking hardware, Big Cloud Fabric leverages software-defined networking (SDN) to make networks intelligent, agile and flexible.

Big switch networks

Big Switch's mission is to disrupt the status quo of networking with order of magnitude improvements in network intelligence, agility, and flexibility by delivering Next-Generation Data Center Networking.

We do this by delivering all the design philosophies of hyperscale networking in a single solution, applicable for data centers of any size.

The Big Cloud Fabric Benefits:

- Physical network intelligence via 'one logical switch' and fabric visibility

- Physical network agility via automation, zerotouch fabric, controller-coordinated upgrading, and faster troubleshooting

- Deployment flexibility via open network hardware and scale-as-you-grow options, for all workloads (physical, VM, and container)

Big Cloud Fabric Overview

Big Cloud Fabric (BCF) is the industry's first data center fabric built using SDN controller software and open networking (white-box or brite-box) hardware switches.

Embracing hyper-scale data center design principles, the BCF solution delivers:

1. Intelligence by simplifying network operations and providing fabric-wide telemetry.

2. Agility by enabling network automation to rapidly deploy application and services.

3. Deployment flexibility powered by a scale-out architecture and open network hardware.

Applications can now take advantage of high east-west bisectional bandwidth, secure multi-tenancy, and workload elasticity natively provided by BCF. Customers benefit from unprecedented application agility due to automation, massive operational simplification due to SDN, and dramatic cost reduction due to hardware (HW)/software (SW) disaggregation.

BCF supports both physical and virtual (multi-hypervisor) workloads and choice of orchestration software.[1] It provides L2 switching, L3 routing, and L4-7 service insertion and chaining while ensuring high bisectional bandwidth. The scalable fabric is fully resilient with no single point of failure and supports headless mode operations.

[1] See specific hypervisor and orchestration support in a later section of this datasheet.
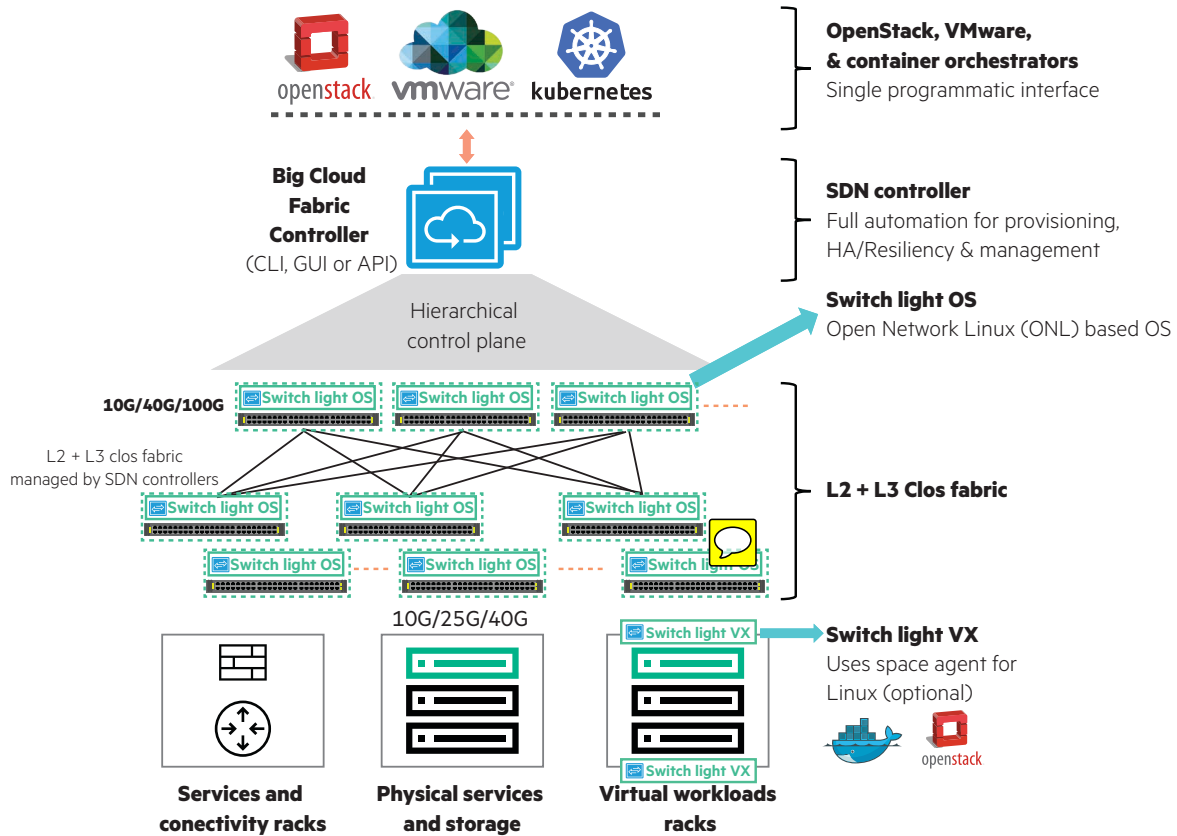
Figure 1: Big Cloud Fabric (Leaf-Spine Clos Architecture)

Architecture: SDN Software Meets open networking Hardware

Software Defined Networking (SDN) fabric architecture refers to a separation of the network's data and control plane, followed by a centralization of the control plane functionality. In practice, it implies that the network's policy plane, management plane and much of 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.

The BCF architecture consists of a physical switching fabric, which is based on a leaf-spine Clos architecture. Optionally, the fabric architecture can be extended to virtual switches residing in the hypervisor. Leaf and spine switches running Switch Light™ Operating System form the individual nodes of this physical fabric. Switch Light Virtual running within the hypervisor extends the fabric to the virtual switches. Intelligence in the fabric is hierarchically placed: most of it in the BCF Controller (where configuration, automation and troubleshooting occur), and some of it off-loaded to Switch Light for resiliency and scale-out.

Big Cloud Fabric System Components

- **Big Cloud Fabric Controller**—a centralized and hierarchically implemented SDN controller available as a High Available (HA) pair of hardware appliances for high availability.

- **Open Networking Leaf and Spine Switches 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 Big Switch hardware compatibility list.

- **Switch Light Operating System**—a light-weight open networking switch OS purpose built for SDN

- **Switch Light VX (optional)**—high-performance user space software agent for KVM-based Open vSwitch (OVS).

- **OpenStack Plugin (optional)**—BSN Neutron plugin or ML2 Driver Mechanism for integration with various distributions of OpenStack

- **VMware vCenter Extension/GUI Plugin (optional)**—built-in network automation and VM Admin visibility for vSphere server virtualization; NSX network virtualization; vSAN storage virtualization, and VMware integrated OpenStack (VIO)

- **Container Plugin (optional)—**BCF plugin for various container orchestrators for container-level network automation and visibility

Deployment solutions

BCF 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 use cases, including:

- VMware SDDC workloads (vSphere, NSX, Virtual SAN and VIO)

- OpenStack including NFV

- Containerized workloads

- Private clouds

- Virtual desktop infrastructure (VDI) workloads

- Big Data/High Performance Computing

- Software Defined Storage (SDS)

The BCF fabric can be designed to support the above listed deployment scenarios using a combination of open networking Ethernet switch options. A few examples are listed in the table shown in Table 1.

Table 1: Example BCF Deployment Scenarios

Example Scenario	Supported Workloads	Leaf Switch Configuration	Spine Switch Configuration
Private/Public Cloud—Typical data center pod deployments	10G	48X10G + 6x40G	32x40G
Cost Optimized Fabric (leverage existing cable infrastructure)	10G	48X10G + 6x40G	48X10G + 6x40G
High performance 40G storage array and 10G workloads (using splitter cables)	10G, 40G	48X10G + 6x40G or 32x40G	32x40G
High Performance Computing/Software Defined Storage/Big Data Analytics	40G	32x40G	32x40G
Dense 10G Compute (using splitter cables)	10G	32x40G	32x40G

Big Cloud Fabric Benefits

Centralized Controller Reduces Management Consoles By Over 60:1

With configuration, automation and most troubleshooting done via the BCF 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 Big 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 BCF 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 Big 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

BCF Controller natively supports integration with various Cloud Management Platforms (CMPs)—VMware (vSphere, NSX Manager, vSAN, & VIO), and OpenStack, and Container orchestrators—through a single programmatic interface. This is tremendously simpler and scalable compared to box-by-box 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.

SDN Fabric Enables Deep Visibility and Resilience for OpenStack Networks

The BCF OpenStack Neutron plugin for L2/L3 networking provides resiliency necessary for production-grade OpenStack deployments—including support for distributed L3 routing and distributes NAT (Floating IP). The BCF Controller acts as the single pane for provisioning, troubleshooting, visibility and analytics of the entire physical and virtual 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.

Network/Security/Audit Workflow Integration

BCF Controller exposes a series of REST APIs used to integrate with application template and audit systems, starting with OpenStack. By integrating network L2/L3/policy provisioning with OpenStack HEAT templates in Horizon GUI, the time to deploy new applications is reduced dramatically as security reviews are done once (on a template) rather than many times (on every application). Connectivity edit and audit functions allow for self service modifications and rapid audit-friendly reporting, ensuring efficient reviews for complex applications that go beyond the basic templates.

Scale-out (Elastic) Fabric

BCF'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 (physical or virtual) 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.
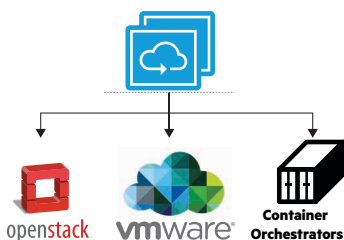

Figure 2: BCF Supports Integration with CMPs

Figure 3: BCF Graphical User Interface (GUI)

DC-grade Resilience

BCF 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 BCF Controller coordinates and orchestrates entire fabric upgrade ensuring minimum fabric down time. These functionalities further enhance fabric resiliency and simplify operations.

Using BCF: A 3-Tier Application Example

BCF 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:

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).
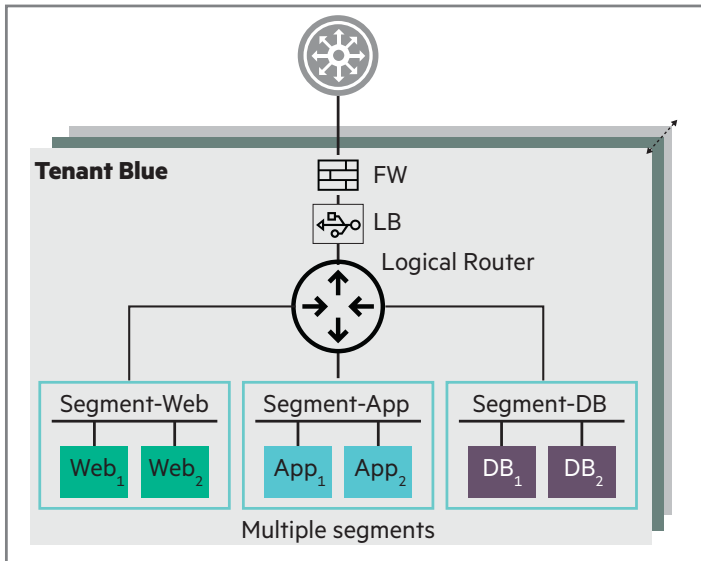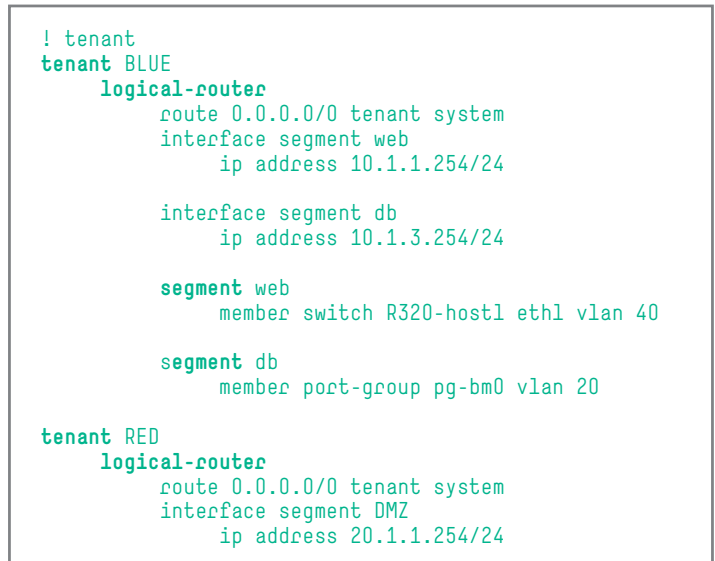
Figure 4: BCF Logical Topology

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

        interface segment db
            ip address 10.1.3.254/24

        segment web
            member switch R320-host1 eth1 vlan 40

        segment db
            member port-group pg-bm0 vlan 20

tenant RED
    logical-router
        route 0.0.0.0/0 tenant system
        interface segment DMZ
            ip address 20.1.1.254/24
```

Figure 5: Application Centric Configuration

Tenant Workflow

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 4, shows various workload nodes of a tenant named "BLUE". Typically, a tenant provisions these workloads using orchestration software such as OpenStack, VMware vSphere, or BCF Controller GUI/ CLI directly. As part of that provisioning workflow, the BCF Controller seamlessly handles enabling 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 4 could be mapped on the pod network. The BCF 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 below, 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 BCF described in earlier sections.

- L3 routing decision is made at the first hop leaf or vleaf 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

Big Cloud Fabric features

Feature	Description/Benefit
Zero Touch Fabric (ZTF)	ZTF enables complete control and management of physical switches within BCF without manually interacting with the switches. It tremendously simplifies day-to-day network operations: • Auto-configuration and auto-upgrade of Switch Light OS, • Automatic topology updates and event notifications based on fabric link state changes. • Auto-scaling of the fabric—adding or removing nodes and/or links within the fabric requires no additional configuration changes on the controller.
Fabric LAG	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.
Fabric Sync	Fabric Sync intelligently synchronizes Controller Information Base (CIB) with fabric node's Forwarding Information Base (FIB) using the OpenFlow protocol. During a topology change, only delta updates are synced across impacted switches. Fabric Sync ensures strong CIB-FIB consistency, as it is the single point of control for maintaining all forwarding and associated policy tables.
Resilient Headless Mode	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.
Centrally-managed Fabric (GUI, CLI & REST APIs)	Big Cloud Fabric Controller provides single-pane-of-glass for entire fabric. • Administrators can configure, manage, debug or troubleshoot, and upgrade the fabric nodes using CLI, GUI, or REST API. • REST APIs, CLI and GUI have application and tenant awareness. 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.
Fabric Analytics	Fabric Analytics is the set of features that provides Advanced Multi-node Troubleshooting, Analytics & Telemetry in the Big Cloud Fabric solution.
API-first Fabric	Big 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 BCF controller for mitigation. The BCF GUI and CLI utilize the underlying REST APIs—hence are by definition consistent and hardened.
Tenant-aware Fabric	Big 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.
Service-aware Fabric	Big Cloud Fabric supports L3 virtual and physical service insertion and service chaining. Services can be shared across tenants or dedicated to a specific tenant.

Big Cloud Fabric features (continued)

Feature	Description/Benefit
L2 Features	• Layer 2 switch ports and VLAN trunks • IEEE 802.1Q VLAN encapsulation • Support for up to 4K VLANs (i.e. 4K Logical Segments) • MAC address based segmentation • BPDU Guard • Storm Control • MLAG (up to 16 ports per LAG) • 3,800 IGMP Groups • IGMP Snooping • Static Multicast Group • Link Layer Discovery Protocol (LLDP) • Link Aggregation Control Protocol (LACP): IEEE 802.1AX • Jumbo frames on all ports (up to 9216 bytes) • VLAN Translation • Primary/Backup Interface • VXLAN Support
L3 Features	• Layer 3 interfaces: Routed ports, Switch Virtual and Interface (SVI) • Multiple IP-Subnet Support per Segment/SVI • Support for up to 46K IPv4 host prefix, 14K IPv6 host prefix (i.e. Endpoints) • Support for 1K Virtual Routing and Forwarding (VRF) entries (i.e. 1K Logical Routers) • 1K Tenants • Static Route, iBGP, eBGP • 68K IPv4 routes, 8K IPv6 routes • Up to 16 ways Equal-Cost Multipathing (ECMP) • 1K Equal-Cost Multipathing (ECMP) groups • 3K flexible ACL entries • Policy-Based Routing • Multicast Routing • ACL: Routed ACL with Layer 3 and 4 options to match ingress ACL • Jumbo frame support (up to 9216 bytes) • DHCP relay • NAT/PAT support
QoS	• Layer 2 IEEE 802.1p (class of service [CoS]) • Source segment or IP DSCP based Classification • Tenant/Segment based classification • DWRR based egress queuing • CoS based marking • PFC and DCBX
High Availability	• Controller HA • Headless mode (fabric forwards traffic in absence of Controller) • Redundant Spine • Redundant Leaf • Redundant Links • Controller cluster with single Virtual IP • Support for redundant out-of-band management switch

Big Cloud Fabric features (continued)

Feature	Description/Benefit
Security	• Ingress ACLs • Layer 3 and 4 ACLs: IPv4, Internet Control Message Protocol (ICMP), Transmission Control Protocol (TCP), User Datagram Protocol (UDP), etc. • ACLs on controller management interface • ACL logging (IPv4 only) • Control Plane Policing (CoPP) or Rate Limiting
OpenStack Integration	• Nova-network & Neutron ML2 Driver mechanism support • Neutron L3 Plugin support (distributed routing, Floating IP, PAT and Security Group visibility) • Auto Host Detection & LAG Formation • OpenStack Horizon Enhancements (Heat Templates, Fabric Reachability Test) • Dynamic Provisioning of the BCF Fabric • Distributed Routing and NAT • Tenant driven OpenStack Router policy configuration grid • LBaaS support (network automation driven through OpenStack)
VMware vCenter Integration	Provides Fabric Automation and Visibility including: • Auto Host Detection & LAG Formation • Auto L2 Network Creation & VM Learning • Network Policy Migration for vMotion/DRS • VM-level Visibility (VMname, vMotion) • VM-to-VM Troubleshooting (Logical & Physical) • L3 configuration via vSphere web-client plugin • Test Path visibility through vCenter
VMware NSX-v Support	Close the overlay/underlay gap for visibility and troubleshooting. Features include: • Auto Host Detection & LAG Formation • Auto Network Creation for VTEP Port-Group & VTEP Discovery • Underlay Troubleshooting – VTEP-to-VTEP connectivity • Underlay Visibility through Fabric Analytics (VM-name, VXLAN ID, Logical Switch)
VMware vSAN Support	• Auto-detection and LAG formation for vSAN node • Auto-creation of vSAN transport network • vSAN cluster communication troubleshooting for unicast and multicast • Simplified Layer 2/Layer 3 multicast deployment for vSAN transport • vSAN Analytics
Container Support[2]	• Container network automation • Container-level visibility • Container-to-container troubleshooting
Multi-Orchestration Support	Support Multiple Virtual PODs (vPODs) on single BCF Fabric
Inter-Pod Connectivity	• L3—Using Static Route and BGP • L2—Dark Fiber • L2—VXLAN (beta) • Hub and Spoke topology (scale)
MIBs	Documented in a separate MIB's document

[2] Container networking support is experimental (with Kubernetes).

Big Cloud Fabric features (continued)

Feature	Description/Benefit
Industry Standards	• IEEE 802.1p: CoS prioritization • IEEE 802.1Q: VLAN tagging • IEEE 802.3: Ethernet • IEEE 802.3ae: 10 Gigabit Ethernet • IEEE 802.3ba: 40 Gigabit Ethernet
Support for Open Networking Ethernet Switch	Support Broadcom Trident-II, Trident-II+ & Tomahawk ASICs for 10G, 25G, 40G and 100G switch from HPE Altoline. The following models are supported: HPE Altoline 6921, Altoline 6941 and Altoline 6960. For the complete list of supported switch and configurations as well as optics/cables, included in the Big Cloud Fabric Hardware Compatibility List (HCL), please **https://www.hpe.com/us/en/networking/data-center.html**.
Fabric Management	• GUI (IPv4/IPv6) • CLI (IPv4/IPv6)—based console to provide detailed out-of-band management • REST API (IPv4/IPv6) Support • Switch management using 10/100/1000-Mbps management through controller • Beacon LED (based on underlying switch) • Configuration synchronization • Configuration save and restore • Secure Shell Version 2 (SSHv2)—IPv4/IPv6 • Username and passwords authentication • TACACS+/RADIUS—IPv4/IPv6 • Control Plane Security (CPSec)—Encrypted communication between Controllers and Physical/Virtual Switches • Syslog (4 servers)—IPv4/IPv6 • SNMP v1 and v2c—IPv4/IPv6 • sFlow® support • SPAN with Policy/ACL • Fabric SPAN with Policy/ACL • Connected device visibility • Ingress and egress packet counters per interface, per segment, and per tenant • Network Time Protocol (NTP)—IPv4/IPv6 • Test Path—Enhanced Troubleshooting & Visibility with logical and physical fabric views (VM<--> vLeaf <--> Leaf <--> Spine <--> Leaf <--> vLeaf <--> VM) • Fabric Analytics including telemetry and enhanced analysis

BCF Controller Appliance Specification

The BCF 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). Physical appliance is also available in NEBS form factor.

BCF Controller—Physical Appliance Specification The BCF controller is available as enterprise-class, 2-sockets, 1U rack-mount physical appliance designed to deliver the right combination of performance, redundancy and value in a dense chassis. It comes in two versions—Standard and Large.

Feature	Technical Specification	
Controller	HWB HP DL360 Gen9 JL553A (Standard)	HWBL HP DL360 Gen9 JL554A (Large)
Processor	Intel® Xeon® Processor E5-2620 v4 20M Cache, 2.10 GHz # of Cores 8 # of Threads 16 Processor Base Frequency 2.10 GHz Max Turbo Frequency 3.00 GHz Cache 20 MB SmartCache Bus Speed 8 GT/s QPI # of QPI Links 2 TDP 85 W VID Voltage Range 0	Intel® Xeon® Processor E5-2650 v4 30M Cache, 2.20 GHz # of Cores 12 # of Threads 24 Processor Base Frequency 2.20 GHz Max Turbo Frequency 2.90 GHz Cache 30 MB SmartCache Bus Speed 9.6 GT/s QPI # of QPI Links 2 TDP 105 W VID Voltage Range 0
Form Factor (H x W x D)	1U Rack Server SFF Drives: 3.44 x 17.54 x 26.75 in (8.73 x 44.55 x 67.94 cm)	
Memory	4 x HPE 16GB (1x16GB) Single Rank x4 DDR4-2400 CAS-17-17-17 Registered Memory Kit	
Hard Drive	2 x HP 1TB 6G SATA 7.2K rpm SFF (2.5-inch) SC Midline 1yr Warranty Hard Drive	
Networking	HPE Ethernet 10Gb 2-port 560SFP+ Adapter HP H240ar FIO Smart HBA	
Power	2 x Hot Plug Power Supplies 500W	
Additional Features	Fan fault tolerance; ECC memory, interactive LCD screen; ENERGY STAR® compliant	
Temperature—Continuous Operation	10°C to 35°C (50°F to 95°F)	
Temperature—Storage	-30°C to 60°C (-40°F to 149°F) with a maximum temperature gradation of 20°C per hour	
Relative Humidity—Continuous	8% to 90% with 24°C (72.5°F) maximum dew point	
Relative Humidity—Storage	5% to 95% at a maximum wet bulb temperature of 38.7°C (101.7°F)	
Altitude—Continuous	3050m (10,000ft)	
Altitude—Storage	9144m (30,000ft)	

Note: For NEBS appliance details please contact the HPE Networking Sales +1-888-269-4073

VM Appliance Specification

The Big Cloud Fabric Controller is available as a Virtual Machine appliance for the following environments (for limited scale production or lab deployment).

Environment	Lab Only	Production
Linux KVM	Ubuntu 14.04	
Vmware ESXi	Version 6.0	
Red Hat RHEL	RHEL 7.2	RHEL 7.2
vCPU	10 vCPU	12 vCPU
vMemory	56 GB of Virtual Memory,	56 GB of Virtual Memory,
HDD	300GB HDD	300GB HDD
vNIC	4 vNIC	4 vNIC

Note: The above table explicitly indicates the Major/Minor/Maintenance versions tested and supported by Big Cloud Fabric. Versions other than the ones listed above will not be supported.

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

Supported workloads & Orchestration systesm

Feature	Technical Specification
Physical Workloads	Bare-metal server workloads
Virtual Workloads	VMware Integration with vSphere 6.0. For OpenStack see table below. Support any virtual workload on BCF P Fabric without integration.
Cloud Orchestration	OpenStack (Neutron ML2 driver, Neutron L3 Plugin) VMware VIO

Open Stack Integration

Hypervisor	OpenStac k – MITa kA (beta)	OpenStac k – Liberty
KVM	Ubuntu 14.04	Ubuntu 14.04 (Mirantis OpenStack—Fuel 8.0)
	CentOS 7.2 (Packstack)	CentOS 7.2 (Packstack)
	RHEL 7.2 (RHOSP 9)	RHEL 7.2 (RHOSP 8)

Learn more at
hpe.com

Make the right purchase decision. Click here to chat with our presales specialists.

f 𝕏 in ✉

Sign up for updates