

A Better Networking Approach for Hyper-Converged Infrastructures



How Big Cloud Fabric Offers Network-as-a-Service for Nutanix HCI Deployments

In January 2019, Big Switch Networks presented their Cloud-First Networking architecture and solution benefits to the Nutanix Technical Champions. This document is in follow up to this presentation, where the technology benefits of Big Cloud Fabric (BCF) over legacy network designs are explained, in both a technology primer section, and in response to several of the questions that were asked during this presentation. The purpose of this document is to inform this community, on an approach to networking, that better aligns with many of the core benefits of hyper-converged infrastructures. As the old saying goes, the chain is only as strong as its weakest link. Networking, if not architected correctly, can mitigate many of the productivity benefits of hyper-converged infrastructures.

Big Cloud Fabric makes networking an invisible extension to hyper-converged infrastructures, while insuring that the core requirements including performance, reachability, and security are fully met. More broadly stated, the network becomes a service where the tenant, storage, and workload networking needs are automatically provisioned with pre-defined policies based on the actions at a cloud layer. With Nutanix, this cloud layer is either through Prism, vCenter, or some other orchestrator. BCF does away with manual configurations, lengthy provisioning wait periods, complex protocols, request tickets, expensive service contracts, and proprietary hardware. All of these aforementioned are artifacts of legacy networking designs that are being force fit (custom scripts) in meeting the needs of cloud computing.

Quick Tech Primer: The Shift from Legacy Branded Offerings to Cloud First

Legacy networking, which follows an operations model that's a quarter century old, boils down to configuration and management box by box. The legacy approach acts as a productivity bottleneck for compute and storage teams. Any manual networking change—a new switch, line card, tenant group, hyper-converged node, or even simple upgrades or security patches—includes wait times that stymie agile development teams. While HCI nodes can easily plug into any legacy data-center-class switch, box-by-box management requires a lot of scripting, proprietary APIs, and highly specialized network staff to orchestrate and automate the network in concert with hyper-converged infrastructures.

The Big Cloud Fabric offering is architecturally different. BCF learns and inherits the configuration states from the cloud orchestration platform. BCF pushes out the inherited states as network configurations to all switch nodes via redundant high-availability controller nodes. The configuration events within Prism drive all the required configuration actions in the BCF physical switch fabric, including moves, adds, and deletions of HCI nodes, tenant pools, storage pools, and virtual machines.

The integration between BCF and the orchestrator (Prism, vCenter, etc.) is easily configured via BCF downloaded plugins, and simple menu driven configuration inputs. Once the information exchange between the BCF controller and the cloud orchestrator is established, everything from that point forward required for auto-provisioning the BCF fabric is auto-learned and then executed.

Big Cloud Fabric offers automates many of the tasks associated with installing new servers, tenants, software upgrades, workloads, and storage nodes. These include the following:

- Leaf switch multi-chassis link aggregation (MLAG) configuration of a newly added HCI node, without any wiring diagrams or port mappings
- Configuration of the HCI topology attributes, including redundancy, failover, load balancing, and quality of service (QoS)
- Dynamic virtual machine (VM) migration configuration
- Auto-creation of the tracking and troubleshooting database to enable tracing all endpoints (virtual or physical) across the BCF topology

BCF controllers manage hundreds of configurations across the entire leaf-spine fabric as one big logical switch. The BCF approach minimizes error-prone CLI configurations. This approach safeguards against misconfiguration, which often create security holes. Further, network specialists can focus their energies on more strategic initiatives as day-to-day change management tasks are fully automated by the BCF controller. And for organizations re-aligning where silo'd teams become combined as one hyper-converged infrastructure team, networking is made easier where the level of competence to deploy, configure and operate the network can be managed by the broader team. This direction aligns with most enterprise infrastructure organizations, where the days of silo's operations, with layers of complexity and bureaucracy, are changing in response to agile application and business needs.

When it comes to business benefits of the Cloud-First approach of BCF, at the top of the list are greatly reduced time for adding HCI nodes, tenant pools, virtual machines, and tuning performance. You can achieve productivity in days, not weeks or months. Your team benefits from the network becoming an extension of your HCI solution, where the network responds to

changes at the speed of a virtual machine with container agility—within minutes, not days. The cloud-first architecture of BCF makes traditional box-by-box networking and manual CLI-driven management a thing of the past. The network is vital, but when it's simplified it becomes transparent, responding at the same velocity as other parts of HCI infrastructure—the speed of VMs.

Frequently Asked Questions

Does BCF create VLANs automatically on the HCI nodes?

This is a good example of how the BCF cloud-first technology works. The Big Cloud Fabric controller automatically learns about the VLANs created within the HCI infrastructure, based on the plug-ins with Prism and vCenter, and automatically configures them within the BCF fabric. This is a zero-touch configuration approach. The tenant community or HCI administrator drives the creation of the VLANs, not the network community. The tenant requests an HCI VLAN on the HCI node from the HCI orchestrator (Prism, vCenter). The BCF Controller learns of the VLAN and automatically pushes it out to the switch interfaces connected to the HCI nodes.

Is the Big Cloud Fabric Controller based on a virtual machine, or some other form factor?

You can configure the x86-based BCF Controller as a virtual machine or as an appliance.

Does BCF Controller have failover and high-availability capability?

Yes. When deployed in a high-availability (HA) configuration, a BCF fabric has no single points of failure—including the data, control, and management planes. Big Switch provides HA reference architecture.

What happens in a catastrophe if all the HA controller nodes fail or lose communication with the fabric switches?

Does the network stop working?

No, the Big Cloud Fabric data plane continues to operate in headless mode. The network continues to forward traffic.

Does the BCF Controller configure anything within the HCI nodes?

In the current design, Nutanix Prism is the single source of truth. We would like to explore new workflows related to networking with the Nutanix user community. [Here is how to provide your input.](#)

Is a command line interface (CLI) available?

Yes. Big Cloud Fabric supplies a traditional, networking friendly CLI on the controller.

What are the alternatives to the CLI interface?

Many customers prefer the graphic user interface, which customers have said is the best GUI in the networking world. Some customers leverage the published APIs in their service-orchestration workflows.

Does Big Cloud Fabric interoperate with Nutanix Flow?

Yes, BCF interoperates with Nutanix Flow. We look forward to partnering with Nutanix Champions to explore advanced Flow-related use cases. [Here's how to provide your input.](#)

Is BCF better suited for HCI brownfields or greenfields?

You can deploy Big Cloud Fabric in either a brownfield or greenfield scenario. To date most BCF deployments are in brownfield data centers where BCF is used as a pod fabric. Communication between the BCF fabric to the existing data center network (either traditional network or another BCF pod fabric) takes place by standards-based L3 routing protocols, such as Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF).

Can BCF extend virtual extensible LAN (VXLAN) out over the WAN?

Yes. BCF supports L2 extensions over L3 network, leveraging VXLAN.

Does Big Switch Networks provide both the hardware switches and the Big Cloud Fabric software components?

Big Switch is a software networking company. We have qualified a variety of third-party open networking switches, both brite box and white box. You can purchase switches from a number of open networking switch suppliers, including Dell EMC and Edgecore Networks.

What type of switches are compatible with the Big Cloud Fabric offering?

Big Switch Networks has an extensive list of qualified standards based, data center class, 1RU and 2RU switches, with port speeds of 1/10/25/40/100 Gbps. The qualified multi-terabit, nonblocking, low-latency Ethernet switches have hot swappable fans and power supplies. These are Google style switches, based on Broadcom merchant silicon. The switches ship with standards-based auto-boot Open Network Install Environment (ONIE)—an open source project Big Switch Networks contributes to. Adding a switch to the BCF fabric is a zero-touch process, including bringing up the switch and installing the network operating system. To obtain a current list of qualified switches, [send us an email message](#). Ask for the Hardware Compatibility List (HCL).

What support procedures are in place between switch vendors and Big Switch Networks?

Big Switch Networks and qualified switch vendors work closely together to ensure a “single throat to choke” support model. We’re proud to say that our customers give our support 4.8 out of 5 stars when rating our technical support. You can see our customer reviews at [Peer Insights](#).

Does BCF configure any virtual network settings in ESXi or AHV?

In the current design, Nutanix Prism or VMware vCenter is the single source of truth. Big Switch Networks would like to explore with the Nutanix user community new operational and troubleshooting workflows for the joint solution. [Here’s how to provide your input](#).

Is BCF more targeted for small HCI deployments such as branches or small offices? Or for larger ones with 20+ HCI nodes?

BCF is optimized for both. At a multi-rack level, you would have redundant top-of-rack switches and dedicated x86-based BCF controllers. Within smaller one-rack or sub-rack HCI densities, you would purchase smaller form factor switches (1/2 width RU) and deploy redundant controllers as virtual machines.

Are there plans to integrate with additional ONIE switch vendors?

Big Switch continues to explore new switch vendor partnerships that align with Nutanix go-to-market motion. We’d appreciate any input in this matter. [Here’s how to make your suggestions](#).

How is the BCF licensed and sold?

We license BCF on a per-switch basis, as a software subscription. And customers need to purchase a subscription for the controller, whether it be virtual machine or appliance based.

What was the Nutanix IT use case to deploy Big Switch’s Big Cloud Fabric?

You can hear the answer directly from Raul Robles, Director of Engineering at Nutanix in a [one-minute video](#).

You mentioned a hands-on product demo?

How do I access it?

Go to labs.bigswitch.com for a virtual tour and a hands-on proof-of-concept experience. To read more Big Cloud Fabric product details, check out the [BCF datasheet](#).



Headquarters

3111 Coronado Drive, Building A
Santa Clara, CA 95054

+1.650.322.6510 TEL
+1.800.653.0565 TOLL FREE

www.bigswitch.com
info@bigswitch.com