



BIG CLOUD FABRIC FOR

Nutanix Enterprise Cloud



BIG CLOUD FABRIC FOR NUTANIX ENTERPRISE CLOUD OVERVIEW	
Introduction.....	3
SOLUTION ARCHITECTURE	
Cloud-First Experience	
Enterprise VPC (E-VPC).....	4
Nutanix Prism driven E-VPC Automation.....	4
Cloud-First Infrastructure	
Open Commodity Hardware.....	4
SDN Based Fabric Controller.....	4
Cloud-First Consistency	
Deploy Anywhere – Operate with One Dashboard.....	5
DEPLOYMENT USE CASES	
Small Deployments: BCF for Nutanix AHV Edge Clouds (1 -3 Racks)	
Overview	6
Solution Requirements.....	6
BCF For Nutanix AHV Edge Cloud Solution & Benefits.....	6
Current example customers for this use case.....	7
Large Deployments: BCF for Nutanix AHV Private Clouds (2 or more Racks)	
Customer case study.....	8
Overview.....	8
Solution Requirements.....	8
BCF For Nutanix Private Cloud Solution & Benefits.....	8
BIG CLOUD FABRIC BEST PRACTICES FOR NUTANIX AHV INTEGRATION	
Provisioning Best Practices	
Connectivity best practices for 1G/10G links.....	10
VLAN configuration best practices for AHV Management Network.....	11
Authentication best practices for BCF Nutanix Prism Integration.....	11
LAG Configuration best practices.....	12
Automation Best Practices	
Visibility-Only.....	13
AHV Host Automation.....	14
E-VPC Automation.....	14
Multi-Tenancy Best Practices	
Single E-VPC Mode.....	15
Multi E-VPC Mode.....	16
Scenario Walkthroughs	17
BIG CLOUD FABRIC SOLUTION BENEFITS	
Automation Benefits	
E-VPC Automation for Nutanix.....	21
Automated AHV Host Detection & MLAG Creations.....	21
Flexible & Customized Automation Options.....	22
Distributed Logical Routing.....	23
Visibility and Fabric Analytics	
Integrated Fabric Analytics.....	25
Troubleshooting	
Troubleshooting logical connectivity problems: Test Path – Logical Trace.....	26
Troubleshooting network traffic related problems: Test Path – Fabric Trace.....	27
Spanning traffic from Nutanix AHV nodes.....	28
Multi-Cluster Automation	
Multi cluster connectivity options for AHV and ESXi Hypervisors.....	29
Business Continuity: Extending Network Connectivity Across Regions.....	30
Single Pane of Glass for networking across clusters/regions using Multi Cloud Director.....	30



Big Cloud Fabric for Nutanix Enterprise Cloud Overview

INTRODUCTION

Public cloud providers have spearheaded innovation in data center architecture to enable a frictionless, self-service experience for application deployment. Organizations are eager to bring the same innovation, operational agility and simplicity into managing on-premises applications. Legacy infrastructure has not been able to meet the evolving needs of enterprise applications because they are manual, complex, siloed and hardware-defined. Enterprises need cloud-like infrastructure that is built upon software-defined principles to eliminate complexity, to enable agile provisioning of applications and to dynamically scale infrastructure with the application needs.

The Industry-leading Nutanix Enterprise Cloud Platform provides one of the easiest ways to bring up your on-premise enterprise cloud. With their AHV hypervisor, Nutanix combines virtualization, compute and storage into a hyper-converged architecture, eliminating the need for having a separate storage area network; The entire infrastructure stack, from disks to VMs is managed with a single pane of glass through Nutanix Prism. Just like how compute and storage are provisioned in public clouds, Nutanix enables provisioning of application workloads with a single click through Prism.

Physical networking, however, has remained challenging with outdated, box-by-box CLI-based approach, which simply cannot be operated in cloud timelines. With traditional solutions, deployment of the network takes weeks. Basic ongoing operations require manual, trouble-ticketing processes and complexity multiply when the network needs to be scaled. As a result, a significant chunk of IT effort is spent is on just scaling and managing infrastructure rather than focusing on applications and business value they provide. Achieving a self-service application model for Nutanix Enterprise Cloud requires an approach where standing up a network is as simple as acquiring a public cloud VPC, ongoing network operations are fully automated, and troubleshooting is extremely simplified.

Big Cloud Fabric from Big Switch Networks, is built from the ground up using cloud design principles to provide the features and flexibility of public cloud networking to the on-prem enterprise clouds. BCF integrates with Nutanix Prism and AHV to enable unprecedented operational velocity, network automation at VM speed and end-to-end network visibility, thus making it the ideal networking solution for Nutanix HCI. This integration simplifies various aspects of the Nutanix AHV cloud infrastructure life cycle: starting from bootstrapping up to application workload provisioning all the way up to ensuring Business Continuity. You will be able to deploy, operate and scale out your Nutanix AHV clusters at the speed your business demands by making physical networking hidden and consumable as an on-demand service with the help of Big Cloud Fabric.



Solution Architecture

Big Cloud Fabric is build based on the idea of bringing agility and operational benefits of public clouds to the on-premise datacenters. As a result, Big Cloud Fabric resembles many of the same experiences & principles public cloud providers incorporate in their networking designs. This cloud-first networking approach Big Cloud Fabric takes to bring public cloud benefits to the on-prem datacenters can be broken down to the following categories: Cloud-First Experience, Cloud-First Infrastructure & Cloud-First Consistency.

CLOUD-FIRST EXPERIENCE

Enterprise VPC (E-VPC)

Virtual Private Cloud or VPC is the primary logical public cloud networking construct today. VPCs allow public cloud users to consume networking resources in a more abstracted and elastic manner without getting exposed to the underlying physical networking constructs. Big Cloud Fabric brings the exact VPC concept to on-prem networks and allows users to consume network resources via Enterprise Virtual Private Clouds (E-VPCs). Rather than confined to traditional physical networking resources and constructs, Enterprise VPCs will enable the network to deliver on-demand network-as-a-service experience.

Nutanix Prism driven E-VPC Automation

With Big Cloud Fabric Nutanix Prism Integration, Nutanix Prism will be able to consume physical networking resources on demand through Enterprise VPCs. The entire networking workflow for the Nutanix AHV hosts starting with the auto discovery of Nutanix AHV hosts to L2 switching & L3 distributed routing is automated once the Nutanix Prism integration is enabled on the Big Cloud Fabric,. Nutanix Prism will be able to drive the networking workflows and consume the physical networking resources on-demand without anyone ever touching the physical network. Network admins will be able to control the amount of freedom granted to the Nutanix admins and make sure the level of automation is customized according to organizational needs.

CLOUD-FIRST INFRASTRUCTURE

Open Commodity Hardware

For quite some time, public clouds operators have been enjoying the cost and operational benefits of open hardware for every aspect of public cloud data center design. Nutanix pioneered the same open hardware value proposition for compute & storage for on-prem clouds. However, when it comes to networking, the industry is locked into vendor-specific solutions. Big Cloud Fabric changes this vendor lock-in by giving the freedom to pick and choose Open Networking hardware from different vendors. This cloud-style Hardware/Software disaggregation is a key differentiator in Big Cloud Fabric to make sure the foundation layer of on-prem datacenters is comparable to the way public-cloud datacenters are built.

SDN Based Fabric Controller

Big Cloud Fabric provides a logical abstraction of the entire physical network through a Software Defined Networking (SDN) based Fabric Controller. As a result, the whole network gets abstracted and managed like One Big Switch. This eliminates the need to login to individual switches to do configurations and regular operations. The need for complex networking protocols and turning on various knobs to fine-tune the network are not needed any more with Big Cloud Fabric.

CLOUD-FIRST CONSISTENCY

Deploy Anywhere - Operate with One Dashboard

With Enterprise VPCs, admins now have a common operational model across both on-prem clouds as well as public clouds to consume & manage networking resources. BCF's API-first principle ensures API consistency for fabric-wide programmability. Similar to the way public clouds lets you consume resources using API models, BCF lets you provision, manage and analyze on-prem private cloud networks through logical API interfaces.

Multi-Cloud Director (MCD) can be used as a single pane of glass to manage Big Cloud Fabric deployments across clouds. Big Cloud Fabric can be run on both Brownfield or Greenfield data center private clouds. BCF is also optimized to run on remote sites to provide Edge Cloud capabilities to organizations to facilitate emerging use cases such as IOT. Big Cloud Fabric public cloud edition provides visibility & analytics to public cloud networking. With MCD's centralized visibility and inventory capabilities, admins can keep a tab of networking from a single management console.



Deployment Use Cases

SMALL DEPLOYMENTS: BCF FOR NUTANIX AHV EDGE CLOUDS (1 -3 RACKS)

Overview

In this type of deployments, a small number of Nutanix Clusters (typically 1 Nutanix cluster with 3 Nutanix AHV nodes) are deployed in remote locations to provide edge cloud capabilities to the organization. Application workloads that are getting spawned in these edge datacenters will cater to various use cases, the most prominent ones being IOT applications or remote or branch offices (ROBO). Organizations typically have a large number of such sites, and centralized manageability of these remote sites becomes very important. Nutanix AHV simplifies the application deployments drastically by providing a hyper-converged stack for both compute & storage. In addition, Prism Central provides centralized manageability. However, networking remains quite challenging in these environments, especially when it comes to centralized site management and automation.

Solution Requirements

- Large number of remote sites, with no specialized onsite IT Staff
- Built in resiliency to withstand drastic environmental conditions
- Plug & Play Deployment & Operations

Typical deployments have around 100-500 locations with no specialized resident IT staff in the remote sites to manage the complete solution. Also, the availability of steady power from the main power distribution lines varies significantly from site to site. A successful end solution should expect power downtimes, and in the event of complete power failure, the cluster should be able to recover automatically. As a result, requirements for the end solution are high resilience with very high degree of automation, and the ability for centralized IT staff to manage all the remote sites through a single pane of glass.

BCF For Nutanix AHV Edge Cloud Solution & Benefits

BCF for Nutanix Edge Cloud solution is developed to provide the following key benefits for the Nutanix Edge Cloud Deployments:

- Fully automated & resilient networking
- Workload visibility & analytics
- Centralized site management via single pane of glass

With BCF, various networking operations are automated, and Nutanix Prism will be able to get on-demand network provisioning for both Layer-2 switching as well as Layer-3 routing. BCF eliminates the need for common manual networking tasks such as VLANs, route creation etc. As a result, no manual intervention is needed on the remote site for cluster operations. Big Cloud Fabric for Nutanix AHV Edge Cloud is specially optimized for automated recovery procedures in case of complete power failures. With Nutanix AHV integration, BCF provides visibility and analytics into workloads that are running on AHV nodes. Once the BCF solution is deployed, Multi-Cloud Director (MCD) is utilized to get the single pane of glass capability to manage and operate all the remote sites from a central location. With the help of the MCD, admins will be able to look at the entire snapshot of their remote sites and perform site-wide operations in a centralized manner.

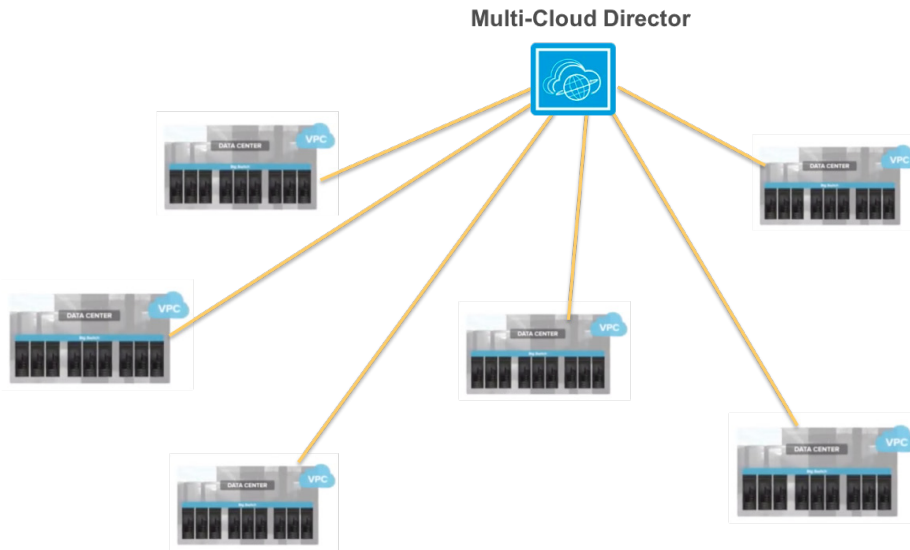


Figure: Multi-Cloud Director (MCD) is used to centrally manage remote sites



Figure: MCD Screen showing status of the remote sites

LARGE DEPLOYMENTS: BCF FOR NUTANIX AHV PRIVATE CLOUDS (2 OR MORE RACKS)

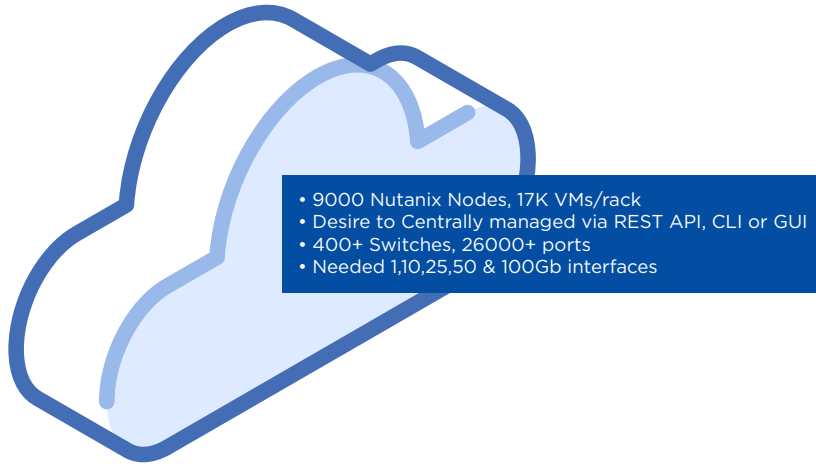
Customer case study

Nutanix IT.

Overview

Nutanix IT provides an IT cloud infrastructure to all the different Nutanix internal teams ranging from Engineering, QA to Field Sales based on Nutanix AHV. The current Nutanix IT Cloud environment has around 9000 Nutanix nodes and spans across multiple data centers locations. Legacy networking was not meeting the demands of the applications and the internal teams. Getting automated multi-tenancy support across various Nutanix AHV clusters and centrally managing networking was highly desired for this private cloud deployments. Legacy physical networking solution that Nutanix IT was using continues to add extra complications and management overheads to these next-generation Nutanix AHV private cloud deployments. Nutanix IT was looking to simplify the networking while adding cloud networking requirements to their physical network.

Deployment Snapshot



Solution Requirements

- Lowering Total Cost of Ownership (TCO)
- On-demand & self-service networking for internal teams
- Automated Multi-Tenancy

Reducing TCO was a primary driver for the redesign. Nutanix IT was spending lots of hours with traditional networking gears to cater to their internal Nutanix team's needs. For example, various Nutanix Engineering teams wanted highly multi-tenant environments to develop and test. Stringent networking requirements such as ability to provide self-serving networking to individual teams & the ability to run overlapping IP address spaces were main drivers.

BCF For Nutanix Private Cloud Solution & Benefits

Big Cloud Fabric was chosen to provide the next generation cloud networking for Nutanix IT Environment. Following were some Big Cloud Fabric key benefits Nutanix-IT highlighted:

- Automation for Nutanix AHV
- Minimal risk/investment protection
- Visibility & Built-in Analytics for AHV workloads
- Simplified operational model, managed by lean team of engineers
- VPC Networking allows cloud-style multi-tenant/on-demand provisioning

Source: <http://www.youtube.com/watch?v=gy7EbN95PRk&t=24m15s>

BCF automation in the Nutanix IT private cloud environment helps cut down admins time drastically. Specifically, the automatic LACP MLAG formation for Nutanix AHV nodes and the ability for Nutanix Prism to consume physical networking as an on-demand service provided immense value to the Nutanix IT. BCF awareness of VM locations, Nutanix AHV nodes, and clusters gave the Nutanix IT team the visibility that they never had with legacy networking solutions.

Another aspect of BCF Enterprise VPC provisioning Nutanix IT immediately found valuable was, to be able to manage the entire environment without any physical moves of cables or AHV nodes. For example, it was a common practice for internal teams to share Nutanix AHV nodes among themselves. Previously with traditional networking, Nutanix IT team had to physically move cables and nodes to provide connectivity, because connectivity was partitioned using traditional VLANs and VRFs. With Big Cloud Fabric, they could provision all those movements at a logical level without taking physical location of the Nutanix AHV hosts into account. All the Nutanix IT admins has to do now is change the automatically discovered AHV node membership from one E-VPC to another E-VPC in the BCF configuration.

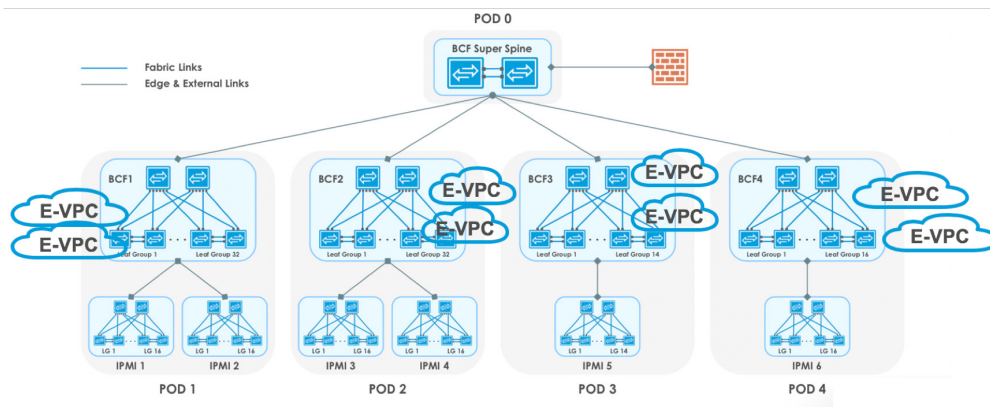


Figure: Example topology diagram showing Nutanix-IT deployment across different datacenter locations

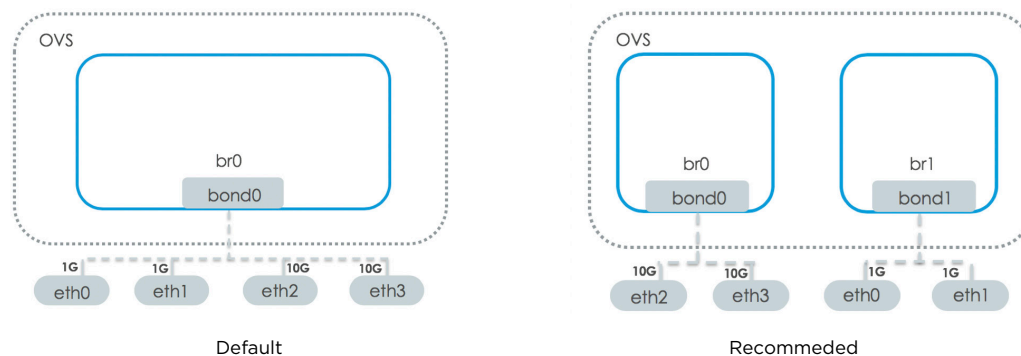


Big Cloud Fabric Best Practices for Nutanix AHV Integration

PROVISIONING BEST PRACTICES

Connectivity best practices for 1G/10G links

By default, all the links in Nutanix AHV nodes are part of the default bridge “br0”. These links include both 1G and 10G links. To prevent suboptimal traffic patterns, it is recommended to remove all the 1G links from the default bridge br0 and attach these links to a newly created bridge.



Alternatively, if you don't want to use 1G links, simply prevent connecting these links to Big Cloud Fabric leaf switches. As a result it is guaranteed that all Nutanix AHV VM traffic always goes through the 10G links.

Steps to remove all the 1G links from bridge, br0:

(Assuming bridge name and bond name are br0 and br0-up, respectively):

```
nutanix@cvm$ allssh 'manage_ovs --bridge_name br0 --interfaces 10g --bond_name br0-up update_uplinks'
```

Optional: To use 1G links if desired, create a new bridge and add the links as follows. Note that it's acceptable to leave the 1g interfaces disconnected instead of adding them to a new bridge.

Create a New Bridge, br1:

```
nutanix@cvm$ allssh 'manage_ovs --bridge_name br1 create_singlebridge
Aggregate 1G interfaces to a separate bond named "br1-up" on the new bridge br1:
nutanix@cvm$ allssh 'manage_ovs --bridge_name br1 --interfaces 1g --bond_name
br1-up update_uplinks
```

Further information around AHV host network management topic can be found at: <https://portal.nutanix.com/#/page/docs/details?targetId=AHV-Admin-Guide-v510:ahv-acr-nw-mgmt-c.html>

VLAN configuration best practices for AHV Management Network

In BCF it is recommended to use a tagged VLAN for AHV management network. Nutanix Foundation Tool can be used to change the VLAN during the AHV cluster setup.

To change the VLAN after the cluster setup log in to individual AHV host and change the VLAN for both AHV host and the CVM.

Complete the following 2 steps using IPMI Out-Of-Band (OOB) management console to change the VLAN on the AHV Host and CVM:

Procedure

Step 1

To change the VLAN on AHV host, SSH in to the AHV host and change the VLAN on the br0 bridge internal port, also named br0.

```
root@ahv# ovs-vsctl set port br0 tag=<vlan>
eg: root@ahv# ovs-vsctl set port br0 tag=4001
```

<https://portal.nutanix.com/#/page/docs/details?targetId=AHV-Admin-Guide-v510:ahv-vlan-setup-ahv-t.html>

Step 2

After changing the VLAN of the AHV host, change the VLAN of the Controller VM (CVM) on that AHV host. Log on to the AHV host and SSH in to the Controller VM (CVM): root@ahv# ssh nutanix@192.168.5.254

Assign the public interface of the CVM to a VLAN:

```
nutanix@cvm$ change_cvm_vlan <vlan-id>
eg: nutanix@cvm$ change_cvm_vlan 4001
```

<https://portal.nutanix.com/#/page/docs/details?targetId=AHV-Admin-Guide-v510:ahv-cvm-assign-to-vlan-t.html>

Follow the same process for rest of the AHV hosts. For details on setting the VLAN configuration back to the defaults, see the Nutanix documentation above.

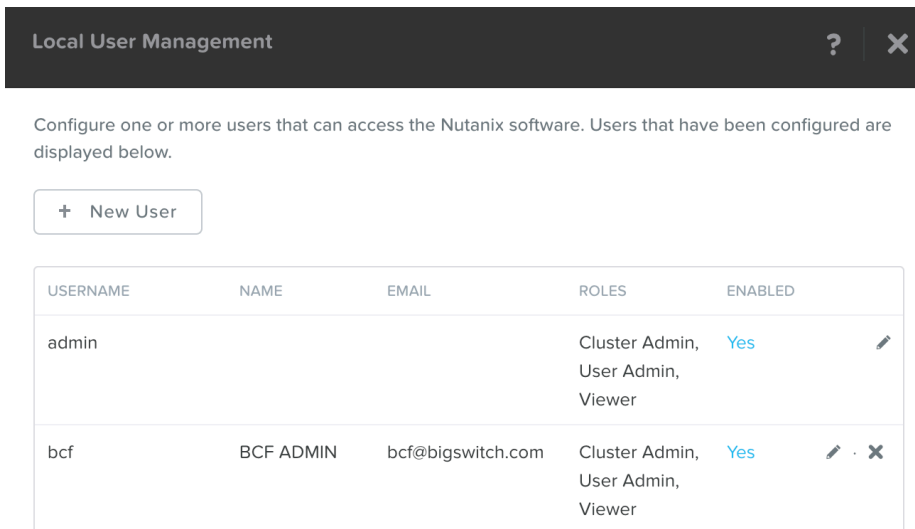
Authentication best practices for BCF Nutanix Prism Integration

Nutanix recommends using a user other than the “admin” user for API communication. To configure a non-admin user for BCF, complete the following steps.

Procedure

Step 1 - In Nutanix Prism Element, select Settings > Local User Management > New User.

Step 2 - Add the “bcf” user with User Admin and Cluster Admin roles.



The screenshot shows the 'Local User Management' interface. At the top, there is a dark header with the title 'Local User Management', a help icon (?), and a close icon (X). Below the header, a text block states: 'Configure one or more users that can access the Nutanix software. Users that have been configured are displayed below.' Underneath this text is a button labeled '+ New User'. Below the button is a table with the following columns: USERNAME, NAME, EMAIL, ROLES, and ENABLED. The table contains two rows of user data.




USERNAME	NAME	EMAIL	ROLES	ENABLED
admin			Cluster Admin, User Admin, Viewer	Yes 
bcf	BCF ADMIN	bcf@bigswitch.com	Cluster Admin, User Admin, Viewer	Yes  

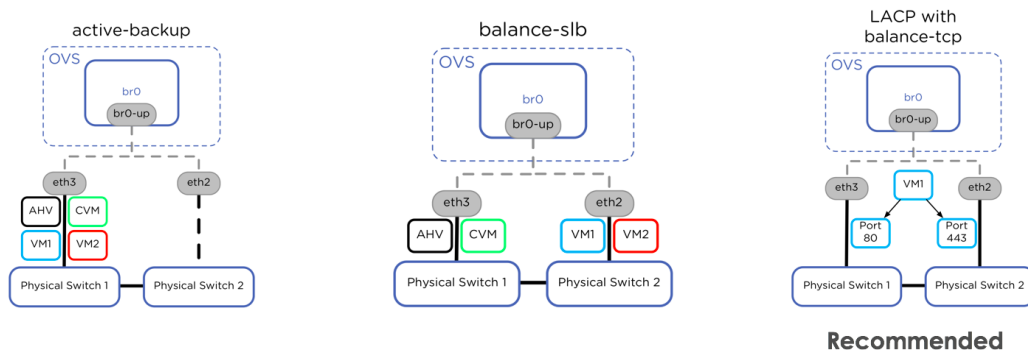
Figure: Nutanix: Local User Management

Step 3 - Use the “bcf” user for the BCF Nutanix Integration.

LAG Configuration best practices

Once the integration is enabled, BCF can automatically detect Nutanix AHV hosts that are connected to the BCF leaf switches with a single link or through multiple links. BCF will automatically figure out the connectivity information through LLDP transmissions and then create appropriate MLAGs (Multi-Chassis LAGs) according to the Prism Configuration.

There are 3 LAG configuration options available in Prism. Active/Backup, Balance-SLB, LACP with Balance-TCP.



Source: <https://next.nutanix.com/blog-40/network-load-balancing-with-acropolis-hypervisor-6463>

In Active-Backup configuration mode only one link will be used as the active link. The other links will always be on the standby mode. Active-Backup is the default load balancing mode in Nutanix AHV. It is always recommended to bootstrap your cluster with this default active-backup mode, and once the Nutanix cluster is running and Integration is enabled on the Big Cloud Fabric gradually move on to the recommended “LACP with Balance-TCP” mode. Note that no manual interventions are needed for migrating Active-Backup to LACP mode.

In Balance-SLB mode one virtual machine will always be pinned down to a particular uplink. From a specific VM’s point of view, only one link will be active. For example, in the above diagram, VM1 will be using eth2 as its active uplink. It will not send or accept any traffic on eth3 interface.

LACP with Balance-TCP provides the best traffic distribution by performing load sharing between both the uplinks. As a result, traffic from a particular VM going to two different destination applications can load share between both the uplinks. It is highly recommended to move to this mode from active-backup once the cluster is bootstrapped and Nutanix Integration is enabled on the BCF. Big Cloud Fabric will automatically handle LAG transitions occurring in Nutanix AHV hosts, including automatically detecting and changing required LACP mode changes in the BCF configuration. No manual user interventions or manual configurations are required on the BCF side. Consult the [Nutanix AHV Networking Best Practices Guide](#) for more information on uplink load balancing.

AUTOMATION BEST PRACTICES

Depending on the individual organizational needs, it is possible to customize automation benefits provided by the Big Cloud Fabric. There are various levels of automation options available in Big Cloud Fabric that may fit many organizational workflows.

Visibility-Only

If you already have customized automation suites in your organization or you don't want to enable automatic configurations in BCF, you can enable just the "Visibility" only mode to get in-depth visibility into the Nutanix AHV clusters without consuming any other automated functionalities. When Visibility-Only mode is enabled, BCF will query Nutanix Prism APIs and populate various visibility views in Big Cloud Fabric. No automated configuration changes such as E-VPC provisioning, segments or distribute logical-routing creations will be performed.

The screenshot displays the Prism-NX1 configuration page in the Big Cloud Fabric (BCF) interface. The page is organized into several sections:

- Info:** A sidebar on the left contains a menu with options: Alerts, VLANs, Graphic (selected), Clusters, Hosts, Virtual Switches, Physical Connections, and Endpoints. Below the menu are "Options & Shortcuts" for content actions (Replace, Append) and placement (Top, Bottom, In order listed above), and "Select All", "Unselect All", and "Restore Defaults" buttons.
- Summary:** A table showing cluster statistics:

Metric	Value
Cluster	1
Hosts	3
Virtual Switches	9
Endpoints	21
- Configuration:** A table of configuration parameters:

Parameter	Value
Name	Prism-NX1
Operational Mode	Maintenance <input checked="" type="checkbox"/> Normal
Configuration Automation Level	Full
Host Name	10.2.17.232
User Name	admin
Last Updated	Today, 7:37:06am GMT
Status	connected
Status Detail	—
CVM Management VLAN	—
CVM EVPC Tenant	—
CVM Segment	—
- Graphic:** A visual overview of the network configuration. It shows three Hosts (Block-A, Block-B, Block-C) and three Virtual Switches (br0, br1, virbr0). Host Block-A is selected, showing its details: 3 virtual switches, 3 physical connections, 12 virtual machines, and 2/2 uplinks connected. The Virtual Switch br0 is also selected, showing its details: 6 endpoints, including CentOS-BP1-4280#tap1, CentOS-BP1-4281#tap3, CentOS-BP1-8652#tap5, CentOS-ST2#tap4, and web-0-1R1101-17592R#.

Figure: Visibility page in BCF showing in depth details of Nutanix Prism network configuration

AHV Host Automation

In addition to providing Visibility, AHV Host Automation mode removes the hassle of manually keeping track of Nutanix AHV links. When enabled BCF will automatically detect AHV links through LLDP and automatically create appropriate interface-groups (or MLAGs) that are corresponding to AHV load balancing settings. BCF is heavily optimized to handle subtle load balancing configuration changes that are happening on Nutanix AHV environments. Manually keeping track of AHV links and MLAGs can lead to misconfigurations as well as it can lead to timing issues which may result in traffic disruptions for Nutanix AHV clusters. Therefore “AHV Host Automation” mode is highly recommended over the “Visibility-Only” automation mode.

E-VPC Automation

In this mode, you will get all the benefits of visibility and AHV host automation plus you'll get to create and map Enterprise VPCs (E-VPC) for Nutanix AHV clusters automatically. When using this mode BCF will automate all the E-VPC configurations including L2 switching as well as L3 routing. When you create VMs and Virtual Networks using Nutanix Prism, BCF will adjust all the necessary network configurations automatically. You have the option of automating networking just for L2 switching, or go all the way up to L3 routing, essentially making networking completely invisible for the Nutanix AHV clusters.

When L3 automation is enabled, whenever the Nutanix admin creates a default gateway in Nutanix AHV Virtual Networks, BCF will detect that and configure Distributed Logical Routing (DLR). As a result, L3 routing will get automated across the fabric for all the Nutanix AHV VLANs. In contrast to traditional networking where you have to manually configure routing protocols & SVI interfaces across different boxes to enable L3 communication, DLR functionality provides 1-Click routing for the entire Nutanix AHV cluster. E-VPC automation also gives you the flexibility of mapping Nutanix AHV clusters to a single or multiple E-VPCs. If you have multi-tenancy requirements in your organization this knob will give automated capabilities to manage multi-tenant networks. Refer to the “Multi-Tenancy Best Practices” section for more details.




Figure: E-VPC automation for Nutanix Prism

MULTI-TENANCY BEST PRACTICES

If you want to create highly multi-tenant environments for your Nutanix AHV clusters, you can use Big Cloud Fabric's flexible E-VPC mapping feature to automate E-VPC creations that suits your organizational multi-tenancy needs. In a high level, the following options are available in BCF:

Single E-VPC Mode

Simplest operational mode and most deployments will fall into this category. The whole Nutanix AHV cluster will be part of this E-VPC and BCF will automate configurations inside this E-VPC. For example, whenever an AHV admin creates a virtual network or a VLAN, BCF will create a network segment inside this E-VPC automatically.




Figure: Single E-VPC mode showing mapping between Nutanix Prism and BCF

Multi E-VPC Mode

If you have multi-tenancy needs in your organization, you can break up the Nutanix AHV cluster to suit your multi-tenancy requirements. BCF provides the ability to break up Nutanix environments by VLANs so that you can assign AHV VLANs to different E-VPCs. The following customization options are available:

By VLAN-Range: In this case, you can map a VLAN range from Nutanix AHV Cluster to an E-VPC. BCF will automatically create network segments for every VLAN in the user-defined VLAN-Range and automatically manage configurations. For example, in the following figure, the QA department will get mapped to its own E-VPC. Whenever a QA user creates a VLAN between 1000-2000, a network segment will automatically get created inside this QA E-VPC.

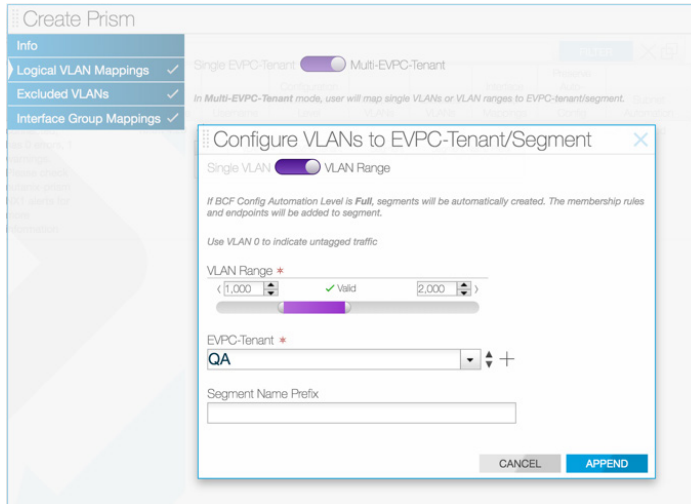


Figure: Nutanix Prism VLAN range mapping to a BCF E-VPC

By Single VLAN: This option can be used to define a more specific VLAN to E-VPC mapping. This single VLAN option can come in handy when Nutanix Prism is not required to drive network segment creation inside BCF automatically. Admin can pre-define an E-VPC and a Segment and mapped that to a specific Nutanix AHV VLAN. In this case, BCF will not automatically manage the E-VPC/Segment creation/deletions. Instead, BCF will only auto-configure VM related configurations inside this network segment.

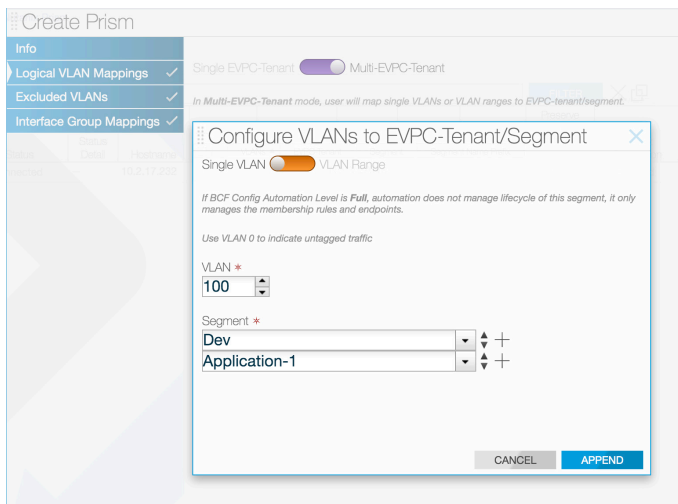


Figure: Single VLAN in Nutanix Prism mapping to a BCF E-VPC & a segment

SCENARIO WALKTHROUGHS

The following scenarios walk through a more practical set of real-world requirements and try to map the requirements to the BCF automation options. This is not an exhaustive list, rather most common scenarios we've seen with customers in production. You might have to mix and match options from different automation modes to fit your organizational needs.

Scenario #1: Converged admins/Network needs to be invisible

Example Admin Profiles:


Networking admins and Cloud admins who manage both Nutanix AHV and physical networking.

Requirements:

Needs plug-&-play networking. Full network automation for Nutanix AHV Clusters. Physical network should be invisible. No Multi-Tenancy requirements.

Recommended Automation Option:

E-VPC Automation provides the best automation option. You want everything to be plug and play and automated, just like in a public cloud. Since you don't have multi-tenancy needs, you can map the entire Nutanix cluster to a single E-VPC. This is the most popular automation mode in Big Cloud Fabric.



Scenario #2: Converged admins/Multi-tenant networking requirements

Example Admin Profiles:

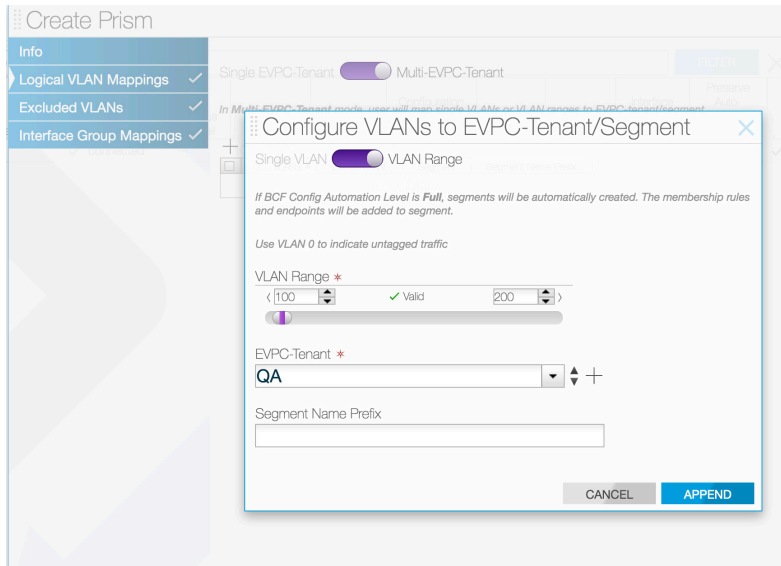
Networking admins and Cloud admins who manage both Nutanix AHV and physical networking.

Requirements:

Wants to provide on-demand Network-As-A-Service to your tenants/customers, and needs multi-tenant plug-&-play networking. Needs full network automation for Nutanix AHV Clusters.

Recommended Automation Option:

E-VPC Automation. To achieve multi-tenancy needs, use the VLAN range option in Multi-EVPC mode to map different VLAN ranges to different tenants. For example, assume you have 2 different tenants called QA and Dev. You can assign VLAN range 100-200 to QA tenant and 300-400 for Dev tenant. Since VLAN range option is used, the entire BCF networking will be completely on-demand from individual tenant's point of view. QA tenant can only consume networking resources when they are mapped between 100-200 VLANs on Prism Virtual Networks.



Scenario #3: Separate Networking & Virtualization teams/Wants to build an on-prem cloud with automated networking

Example Admin Profiles:

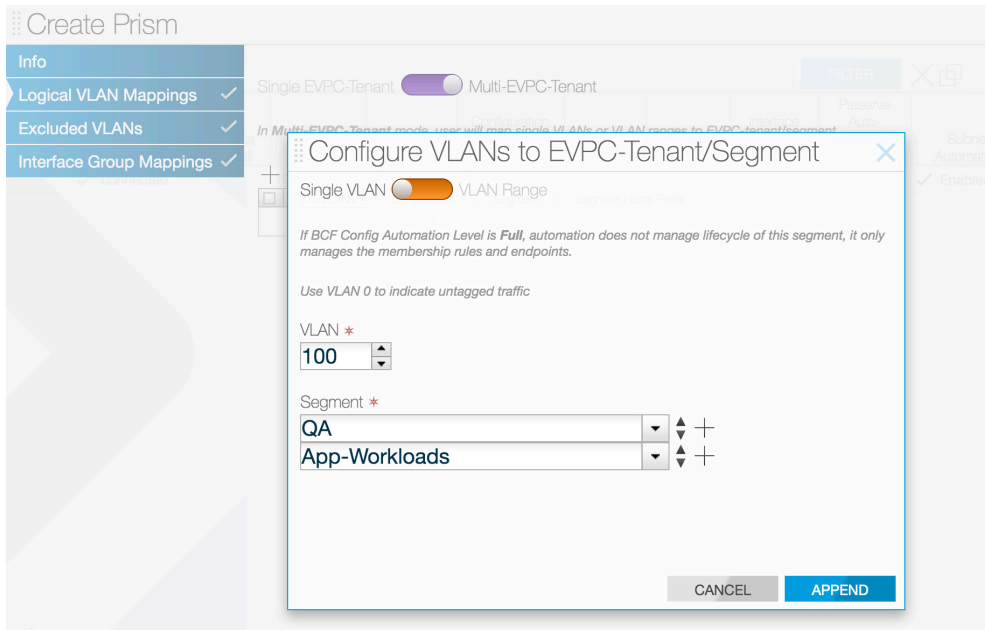
Networking admins. Organization has a separate virtualization team to take care of Nutanix AHV clusters.

Requirements:

Your current operational model is to create VLANs in your network depending on the virtualization team's needs, and once provisioned, hand over that VLAN to the virtualization team. Virtualization team can spin up & down workloads in that VLAN as they wish. Anything outside of that VLAN, virtualization team has to file a ticket to get it approved from your team. You want to move on to a next-gen cloud network, but you still want to maintain the current operational models and organizational best practices.

Recommended Automation Option:

E-VPC Automation (with or without DLR: Distributed Logical Routing, depending on your needs). Regardless of your multi-tenancy needs, in this case, the recommendation is to use Multi-EVPC model option with "Single VLAN" mappings. This way you can provision VLANs, just like the way you would operate today while maintaining full control of network segmentation. Nutanix Prism will not be able to automatically create network segments. However, BCF will help you to automate all the mundane tasks such as MLAG provisioning and Nutanix AHV VM membership provisioning.



Scenario #4: Separate Networking & Virtualization teams/Automation suits are already in place, Needs a next-gen network

Example Admin Profiles:

Networking admins. Organization has a separate virtualization team to take care of Nutanix AHV clusters.

Requirements:

You have automation suites already embedded in your organizational workflows. However, you lack any visibility into the Nutanix AHV environments, and you would like to operate your networks in more cloud-native ways compared to the traditional box by box approaches.

Recommended Automation Option:

AHV Host Automation. This mode will provide you comprehensive visibility into the Nutanix AHV environments, without any config automation. To get rid of manual workflows and to remove yourself from physical configurations of AHV nodes use the “AHV Host Automation” option. BCF will automatically figure out where AHV nodes are connected to the fabric and automate MLAG configurations according to AHV node settings. Once enabled, you don’t have to worry about AHV server moves or any configuration changes with regards to physical server movements; BCF will monitor the AHV node connectivity information and automatically move MLAG formations. As a result, when your server team changes the location of an AHV server, you will not have to perform any network configuration changes on your side.



Big Cloud Fabric Solution Benefits

AUTOMATION BENEFITS

E-VPC Automation for Nutanix

The Nutanix Prism integration allows Big Cloud Fabric to automatically orchestrate Enterprise Virtual Private Cloud (E-VPC) formations for Nutanix Environments. Complete life cycle management of a Nutanix cluster including the cluster bootstrapping is supported with the E-VPC automation. You can plug in Nutanix nodes anywhere in the Big Cloud Fabric without worrying about the exact switch/port locations and get the bootstrapping process established through the Nutanix Foundation tool. Once the cluster is bootstrapped, admins can create the Nutanix Prism Integration in BCF. As a result, Big Cloud Fabric will connect to Prism and gather all the networking related information about the Nutanix AHV clusters. Then BCF will start monitoring the AHV clusters for all ongoing networking related changes such as virtual network creations and deletions, VM creations and deletions, VM migrations, host additions etc. Depending on the desired intent, BCF takes automatic actions to adjust the entire network according to the Prism configurations.

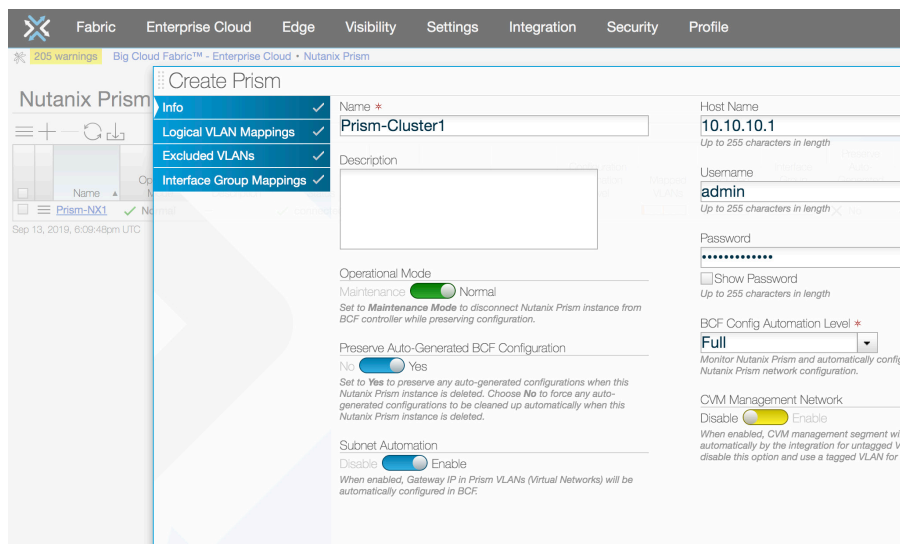


Figure: Configuring Nutanix Prism Integration

Automated AHV Host Detection & MLAG Creations

One of the advantages Big Cloud Fabric brings in with the E-VPC automation is to detect AHV nodes across the fabric automatically and create matching MultiChassis Link Aggregation (MLAGs) configurations according to the Nutanix AHV node settings. As a result, regardless of the physical location of a Nutanix AHV host, it will be mapped to its own E-VPC and will be able to communicate with the rest of the AHV hosts. In the event, if a user moves an AHV host to a different location, Big Cloud Fabric will detect the movement changes through the Link Layer Discovery Protocol (LLDP) and automatically adjust the connectivity mappings without any additional EVPC re-configurations or prompts. Contrast to traditional networking Nutanix admins don't need to worry about network reconfigurations during VM migrations or AHV host movements.

Interface Groups

Name	Description	Leaf Group	State	Admin Status	Mode
nutanix-Prism-NX1-HQCluster-Block-A-br0	Host Block-A cluster HQCluster br0	R2	Up	Up	LACP Fallback Individual

Switch	Switch MAC	Interface Name	Description	Status	Spine Switch	Leaf Switch
R2L1	70:72:cf:ea:1c:e3	ethernet7	-	Up	-	✓
R2L2	70:72:cf:ea:19:b5	ethernet7	-	Up	-	✓

Figure: Automatic formation of MLAGs according to Nutanix Prism Configurations

Flexible & Customized Automation Options

BCF E-VPC automation allows different automation flexibility levels to fit enterprise organizational needs. In most restrictive manner admins can configure Nutanix Prism Integration in BCF to provide visibility into Nutanix AHV clusters without subscribing to any automatic network configuration changes. On the other hand, admins can configure Prism Integration in BCF to provide full fledge benefits, including automated E-VPC configurations for a Nutanix AHV cluster to support multi-tenancy use cases within organizations. For example, IT department serving multiple internal organizations (e.g.: Finance/HR, etc.) can automatically provision multiple E-VPCs for each of their customers. Following is a breakdown of the automation flexibility provided in the Big Cloud Fabric.

AUTOMATION LEVEL	BENEFITS
Visibility Only	Visibility information will be populated in BCF, no automated configuration changes
AHV Host Automation	Visibility + Automatic detection of AHV hosts and Automatic MLAG provisioning
E-VPC Automation	Visibility+ AHV Host Automation+ E-VPC automation to automatically create E-VPCs for AHV clusters Automatically sync L3 default gateway in Prism virtual networks to BCF Distributed Logical Routing

Distributed Logical Routing

Another key benefit of the E-VPC orchestration is the ability to automatically sync the default gateway of each Virtual Network into the BCF Distributed Logical Routing. When an admin maps AHV cluster to an E-VPC, for each Virtual Network (or a VLAN) inside AHV, Big Cloud Fabric will automatically create a network segment inside that E-VPC. When Prism admin creates a default gateway for that virtual network via Nutanix Prism that action can get synced over to the Big Cloud Fabric. As a result, BCF creates distributed logical routing for that virtual network inside the E-VPC. This distributed logical routing happens in an abstracted layer rather than at a physical level as commonly seen with traditional networking. As a result of this automatic sync, the default gateway for each AHV virtual network will be accessible from any physical switch port throughout the fabric in a distributed and abstracted manner without confining to physical limits.

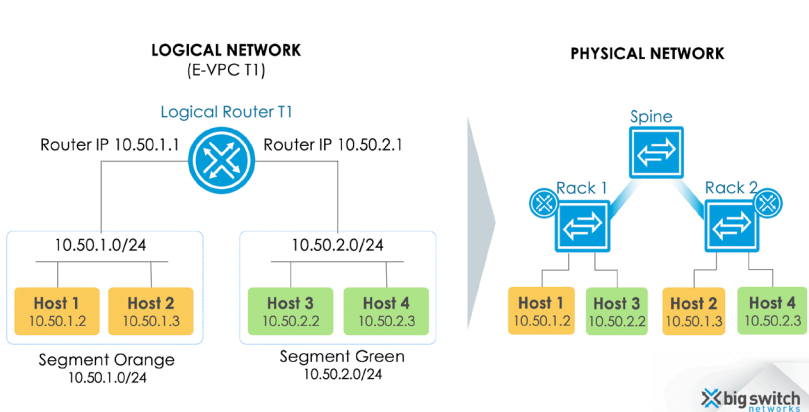


Figure: Distributed Logical Routing for AHV Virtual Networks

VISIBILITY AND FABRIC ANALYTICS

Big Cloud Fabric Prism Integration aspires to bring both Network admins and Nutanix AHV admins to the same page by providing comprehensive visibility into Nutanix AHV virtual networking. BCF will correlate visibility information learned from Nutanix AHV to connectivity information that are learned and configured on the fabric, and it will highlight errors/warnings and provide best practices.

Visibility information such as how Nutanix AHV hosts are connected to the fabric, how MLAGs are formed across the fabric, how different VMs are configured and connected will be presented.

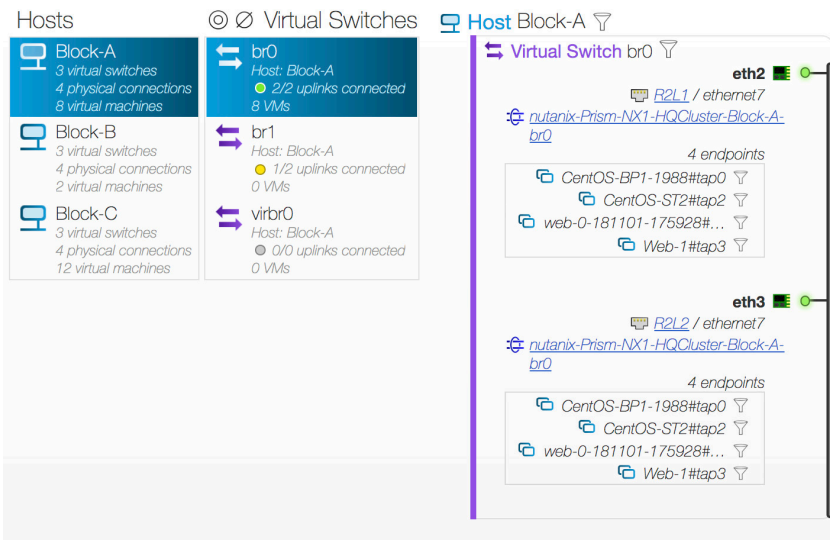


Figure: Comprehensive visibility for Nutanix AHV virtual networks

Admins will be able to navigate across different visibility information targeted by different use cases or problem statements. Following screenshot shows the Nutanix AHV virtual switch specific view which essentially allows network admins to drill down to a particular virtual switch inside the Nutanix AHV host and figure out whether there are any issues from virtual to physical connectivity point of view.

Virtual Switches

Name	Inventory Host Name	Host DNS Name	Status	BCF Switch Inte
br0	Block-A	-	-	

Virtual Switch br0

eth2 ●

R2L1 / ethernet7

[nutanix-Prism-NX1-HQCluster-Block-A-br0](#)

0 endpoints

|

eth3 ●

R2L2 / ethernet7

[nutanix-Prism-NX1-HQCluster-Block-A-br0](#)

0 endpoints

Physical Connections

NIC Name	Connected Interface	BCF Interface Group	Status
eth3	R2L2 / ethernet7	nutanix-Prism-NX1-HQCluster-Block-A-br0	✓ Connected
eth2	R2L1 / ethernet7	nutanix-Prism-NX1-HQCluster-Block-A-br0	✓ Connected

Sep 13, 2019, 6:37:16pm UTC
Show: [10](#) [25](#) [100](#) (1 - 2 / 2)

Figure: Nutanix AHV virtual switch view showing connectivity information

All network related information about endpoints will be aggregated into the Endpoint table, and admins will be able to search and sort through for specific Virtual Machines quickly.

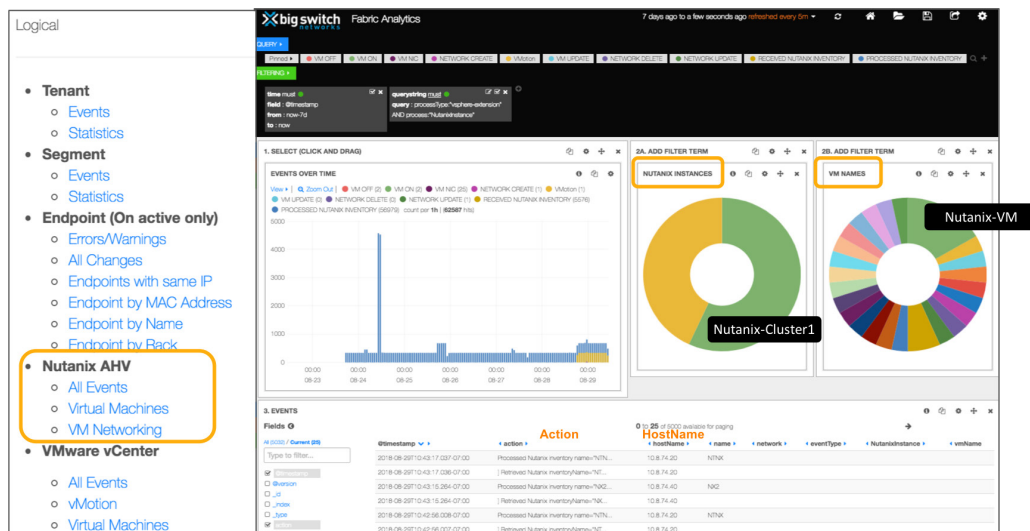
Endpoints

Machine Name	VM Interface Name	Power State	Active in BCF	Segment	EVPC Tenant	Status	VLAN	Inventory Host Name	Physical Connection
App-1	tap2	Powered On	No	Nutanix-VPC-2001	Nutanix-VPC	---	2001	Block-C	eth3 - R112 / ethernet9 eth2 - R112 / ethernet9
CentOS-BP1-1988	tap0	Powered On	Yes	vAcol-VPC-149	vAcol-VPC	---	149	Block-A	eth3 - R212 / ethernet9 eth2 - R212 / ethernet7
CentOS-ST1	tap4	Powered On	No	vAcol-VPC-150	vAcol-VPC	---	150	Block-C	eth3 - R112 / ethernet9 eth2 - R112 / ethernet9
CentOS-ST2	tap2	Powered On	No	vAcol-VPC-150	vAcol-VPC	---	150	Block-A	eth3 - R112 / ethernet9 eth2 - R112 / ethernet9
NTNX-Block-A-CVM	vnet0	Powered On	Yes	Nutanix	Infrastructure	---	---	Block-A	eth3 - R212 / ethernet7 eth2 - R212 / ethernet7
NTNX-Block-A-CVM	vnet1	Powered On	No	---	---	[WARNING] Physical network not configured	---	Block-A	---
NTNX-Block-B-CVM	vnet0	Powered On	Yes	Nutanix	Infrastructure	---	---	Block-B	eth3 - R212 / ethernet8 eth2 - R212 / ethernet8
NTNX-Block-B-CVM	vnet1	Powered On	No	---	---	[WARNING] Physical network not configured	---	Block-B	---
NTNX-Block-C-CVM	vnet0	Powered On	Yes	Nutanix	Infrastructure	---	---	Block-C	eth3 - R112 / ethernet9 eth2 - R112 / ethernet9

Figure: Endpoint view showing networking information for each AHV VM

Integrated Fabric Analytics

Big Cloud Fabric comes with a built-in Fabric Analytics module where all the events related information is automatically archived and analyzed in the background. Fabric Analytics comes with integrated dashboards for Nutanix AHV. All the Nutanix AHV networking information will be archived and presented to the user using different dashboards as shown below. Admins will be able to zoom in to various AHV networks, Virtual Machine events and search and sort through relevant networking information that occurred in a specific timespan.



TROUBLESHOOTING

Troubleshooting logical connectivity problems: Test Path - Logical Trace

Big Cloud Fabric provides with various built-in troubleshooting capabilities for Nutanix AHV environments. “Test Path - Logical Trace” gives admins the ability to run a logical simulation between any two Nutanix AHV VMS (or from AHV VM to any other endpoint such as a container/Bare Metal Host etc). This logical trace will examine all the relevant configuration parameters such as security policies, routing policies related to the source and destination endpoints, and report the trace result back to the user. If the traffic flow is getting blocked due to a particular policy rule in one of the security policies, the logical trace will be able to pinpoint the exact policy rule that’s causing the traffic drop in the fabric.

Logical Trace

Reverse Results Forward Results

> Configuration

▽ Summary

Forward Result: Forwarded

Reverse Result: Forwarded

▽ Logical

Hop Index: 1

Hop Type: endpoint

Endpoint: IP Address: 192.168.150.58
MAC Address: 50:6b:8d:8a:b7:47
Name: nutanix-vm-50-6b-8d-8a-b7-47
Segment: [vApp4-VPC-150](#)
Tenant: [vApp4-VPC](#)

Hop: nutanix-vm-50-6b-8d-8a-b7-47 tenant vApp4-VPC segment vApp4-VPC-150

Policy: —

	Hop Index	Endpoint	Segment Interface	Tenant Router
Source	1	nutanix-vm-50-6b-8d-8a-b7-47 tenant vApp4-VPC segment vApp4-VPC-150		
	→ 2		vApp4-VPC-150	vApp4-VPC
	2			
	← 2		vApp4-VPC-145	
Destination	3	nutanix-vm-50-6b-8d-8a-b7-47 tenant vApp4-VPC segment vApp4-VPC-145		

Figure: Logical Trace showing logical connectivity for Nutanix VMs in Big Cloud Fabric

Troubleshooting network traffic related problems: Test Path - Fabric Trace

Fabric Trace for test path can help to track down the exact hop by hop traffic flow from Nutanix AHV environments. Admins can select the source and destination endpoints using metadata information that are imported to Big Cloud Fabric as a result of the Prism integration, and run a fabric trace to track down specific traffic flow. Upon executing fabric trace, all the switches in the fabric will report the particular traffic flow information back to the BCF controller, and it will present the aggregated report back to the user.



Figure: Fabric Trace showing hop-by-hop physical connectivity for Nutanix VMs in Big Cloud Fabric

Spanning traffic from Nutanix AHV nodes

BCF Fabric SPAN feature allows admins to monitor and SPAN traffic streams out from Nutanix AHV clusters without knowing the exact physical parameters related to the traffic stream. For example, an admin can instruct the BCF controller to SPAN traffic from a particular AHV virtual network with a couple of clicks. In traditional networking, the equivalent task would require not only intimate knowledge about where exactly traffic would be originating but also coordination and configuration amongst many individual network devices and ports. BCF simplifies these SPAN troubleshooting workflows greatly by allowing admins to define the targeted traffic flow parameters using numerous Nutanix AHV metadata information.

The screenshot shows a 'SPAN' configuration interface with a 'Create Filter' dialog box open. The dialog box contains the following fields and options:

- Sequence ***: 1
- 1 - 100**
- Source**: EVPN-Tenant/Segment (Selected: Nutanix-VPC)
- Ethertype**: Select By Name (Selected: Decimal)
- IP Protocol**: Select By Name (Selected: Dec)
- IP DSCP**: [Dropdown]
- IP ECN**: [Dropdown]
- Source IP**: IPv4 Address / Prefix (Subnet Mask)
- Destination IP**: IPv4 Address / Prefix (Subnet Mask)
- Source Port**: [Dropdown]
- Destination Port**: [Dropdown]

Buttons: CANCEL, APPEND

Figure: SPAN Fabric configuration

MULTI-CLUSTER AUTOMATION

Multi cluster connectivity options for AHV and ESXi Hypervisors

Big Cloud Fabric can greatly simplify connectivity options between Nutanix AHV clusters and Nutanix nodes running ESXi hypervisor. With a couple of click process you can map your ESXi cluster to one E-VPC in Big Cloud Fabric and AHV cluster to another E-VPC. As a result, if you are migrating your current ESXi infrastructure to an AHV infrastructure, Big Cloud Fabric can reduce the networking burden entirely by automating many of the associated connectivity workflows. For example, depending on your enterprise needs, you can choose to map both ESXi cluster and an AHV cluster to the same E-VPC thus making sure connectivity is seamlessly maintained across both clusters during the migration period. On the other hand, if you wish to only allow a specific set of connectivity from your current ESXi environment to AHV environments you have the flexibility to pick and choose and customize E-VPC mappings depending on your needs.

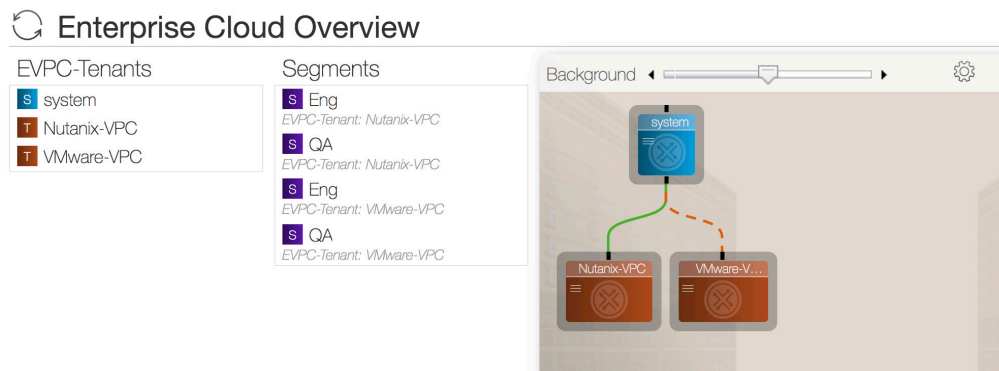


Figure: Connectivity details between Nutanix & VMware VPCs in Big Cloud Fabric

Business Continuity: Extending Network Connectivity Across Regions

Maintaining connectivity across Nutanix AHV clusters in different regions may be crucial to maintaining business continuity. If you operate multiple Nutanix clusters in different datacenters, you can extend the connectivity between these AHV clusters with the help of Big Cloud Fabric. BCF provides two different flavors of connectivity options to connect across regions: Directly connecting BCF pods (e.g.: connecting with dark fiber), or connecting BCF pods using VXLAN over a Layer3 network. Depending on the needs admins can pick and choose between these two flavors to establish inter-region connectivity across geographically dispersed Nutanix AHV clusters.

Single Pane of Glass for networking across clusters/regions using Multi Cloud Director

Multi Cloud Director (MCD) from Big Switch allows you to manage multiple Big Cloud Fabric pods from a single pane of glass. MCD federates data from all BCF deployments, with a centralized database, centralized GUI, and centralized search intelligence. The rich GUI offers at-a-glance dashboards. Admins can quickly visualize data across multiple fabrics and take immediate action in response to anomalies, alerts, and on-screen reports.



info@nutanix.com | www.nutanix.com | [@nutanix](https://twitter.com/nutanix)