



January 8, 2016

53-1003933-02

Nutanix Hyperconverged Appliance with the Brocade VDX ToR Switch

Deployment Guide

© 2016 Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, Brocade Assurance, the B-wing symbol, ClearLink, DCX, Fabric OS, HyperEdge, ICX, MLX, MyBrocade, OpenScript, VCS, VDX, Vplane, and Vyatta are registered trademarks, and Fabric Vision is a trademark of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned may be trademarks of others.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

The authors and Brocade Communications Systems, Inc. assume no liability or responsibility to any person or entity with respect to the accuracy of this document or any loss, cost, liability, or damages arising from the information contained herein or the computer programs that accompany it.

The product described by this document may contain open source software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Brocade Proprietary and Confidential. The information contained herein is the property of Brocade. The possessor of this document agrees to maintain it in confidence, not to reproduce, copy, reveal, or publish in whole or part.

Contents

- List of Figures 4
- Introduction 5
- Brocade VDX 6740 Top-of-Rack Switch 5
- Nutanix Hyperconverged Appliance 5
- Initial Nutanix Installation 6
- Nutanix Deployment with the Brocade VDX 6740 7
 - Brocade VDX Setup and Configuration 7
 - Firmware Upgrade Procedure 8
 - Configuring the Two Brocade VDX Switches in a Single Fabric in Logical Chassis Mode 9
 - Configuring Brocade VDX Appliance-Facing Ports 11
 - Configuring VDX In-band Management 12
 - Nutanix Setup and Configuration 12
 - Verifying Connectivity Between Nutanix and Brocade VDX 16
- Conclusion 17
- References 17

List of Figures

Figure 1	Brocade VDX 6740 1U ToR Switch	5
Figure 2	Sample Nutanix Appliance.....	5
Figure 3	Nutanix Software Architecture.....	6
Figure 4	Example of a 3-Node Nutanix Deployment.....	7
Figure 5	Nutanix Deployment with Brocade ToR Switches	8
Figure 6	Brocade VDX Firmware Download from USB.....	9
Figure 7	Brocade VDX Firmware Download from FTP Server.....	9
Figure 8	Brocade VDX—Enabling Logical Chassis Mode on VDX Switch 1.....	10
Figure 9	Brocade VDX—Enabling Logical Chassis Mode on VDX Switch 2.....	10
Figure 10	Brocade VDX—Show Output Displaying Logical Chassis Mode and Firmware Version	11
Figure 11	Brocade VDX 6740 Configuration on Appliance-Facing Ports	12
Figure 12	Brocade VDX 6740 Configuration for In-band Management.....	12
Figure 13	Nutanix Prism GUI Login.....	13
Figure 14	Nutanix Prism GUI Home Screen	13
Figure 15	Creating a Container from the Nutanix Prism GUI	14
Figure 16	Default and User-Created Containers in the Prism GUI.....	14
Figure 17	Network Configuration in the Nutanix Prism.....	15
Figure 18	Creating a VLAN ID.....	15
Figure 19	Show MAC Address Table on the Brocade VDX	16
Figure 20	Nutanix Node Details from Prism.....	16
Figure 21	IPMIP MAC Address Details from the IPMI Screen	17

Introduction

This document provides the details for a basic deployment of the Nutanix hyperconverged storage appliance with Brocade VDX top-of-rack switches.

Brocade VDX 6740 Top-of-Rack Switch

The Brocade® VDX® 6740 (Figure 1) offers 48 10-GbE SFP+ ports and 4 40-GbE QSFP+ ports. Each 40-GbE port can be broken out into four independent 10-GbE SFP+ ports, providing an additional 16 10-GbE SFP+ ports. In addition, the switch features low power consumption, consuming 2 watts per 10-GbE port. The switch comes in a 1U footprint.



Figure 1 Brocade VDX 6740 1U ToR Switch

Nutanix Hyperconverged Appliance

The Nutanix appliance is a converged storage + compute solution that can combine multiple scale-out nodes to create a distributed platform for virtualization, also known as a virtual computing platform. The Nutanix solution is a bundled hardware and software appliance delivered in a variety of form factors from 1 node in a 1U chassis up to 4 nodes in a 2U chassis (called a block). Multiple nodes are combined to form a cluster of compute and storage.



Figure 2 Sample Nutanix Appliance

Each node runs an industry-standard hypervisor (ESXi, AHV, or Hyper-V currently) and the Nutanix Controller VM (CVM). The Nutanix CVM runs the Nutanix software and serves all of the I/O operations for the hypervisor and all VMs running on that host. For the Nutanix units running VMware vSphere, the SCSI controller, which manages the SSD and HDD devices, is directly passed to the CVM leveraging VM-Direct Path (Intel VT-d). In the case of Hyper-V, the storage devices are passed through to the CVM. The software architecture is shown in Figure 3.

Nutanix Prism is a management and monitoring console tightly integrated with the Nutanix cluster. Nutanix Prism manages one or more Nutanix clusters.

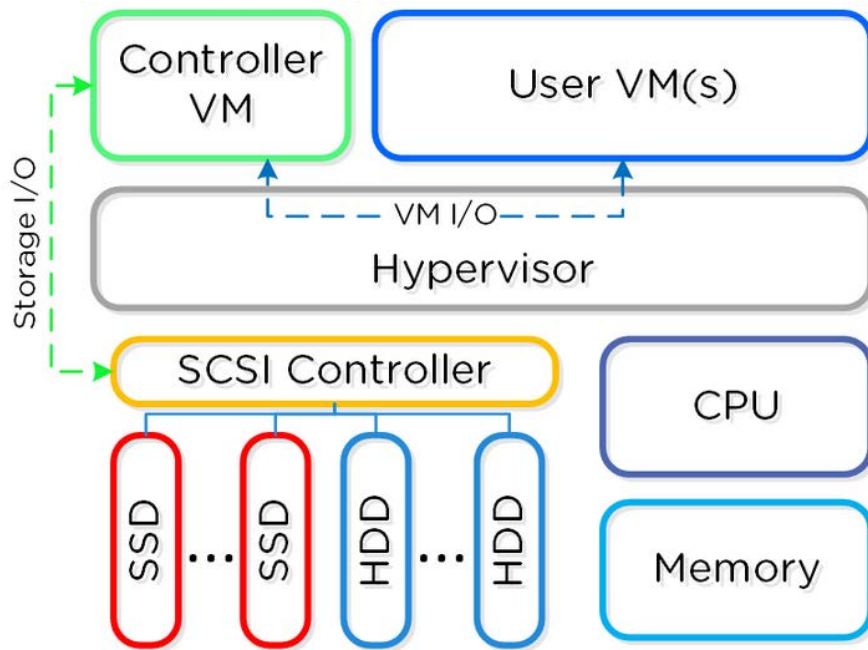


Figure 3 Nutanix Software Architecture

Initial Nutanix Installation

Nutanix installs the Acropolis Base Software Controller VM and the AHV hypervisor at the factory before shipping each node to a customer. To use a different hypervisor (ESXi or Hyper-V), nodes must be re-imaged in the field. The *Nutanix Field Installation Guide* provides step-by-step instructions on how to re-image nodes (install a hypervisor and then the Acropolis Base Software Controller VM) after they have been physically installed at a site.

https://portal.nutanix.com/#/page/docs/details?targetId=Field_Installation_Guide-v3_0:v3_overview_foundation_c.html

Note: Only Nutanix sales engineers, support engineers, and partners are authorized to perform a field installation. The basic steps are outlined below for quick reference. Detailed content can be found in the link above.

A field installation includes the following steps^[3]:

Imaging a Cluster (standard method):

1. Set up the installation environment as follows:
 - a. Download the Foundation Applet (a multinode installation tool) and the Nutanix Installer Package files to a workstation. Also, acquire an ESXi installer from the customer and download it to the workstation.
 - b. Connect the Ethernet ports on the nodes to a switch.
2. Open the Foundation Web GUI on the workstation and configure the following:
 - a. Enter the hypervisor and IPMI address and credential information.
 - b. Select the Nutanix Installer Package and the hypervisor ISO image files to use.
 - c. Start the imaging process and monitor the progress.

Imaging is performed from a workstation with access to the IPMI interfaces of the nodes in the cluster. Imaging a cluster in the field requires first installing certain tools on the workstation and then setting the environment to run those tools. This requires two preparation tasks:

1. Prepare the workstation:

Preparing the workstation can be done onsite or offsite at any time prior to installation. This preparation includes downloading ISO images and running the node discovery application.

2. Set up the network:

The nodes and workstation must have network access to each other through a switch at the site.

- a. Connect the first 1-GbE network interface of each node to a 1-GbE Ethernet switch. The IPMI LAN interfaces of the nodes must be in failover mode (factory default setting).
- b. Connect the installation workstation to the same 1-GbE switch as the nodes. The first 1-GbE network interface on a node is a shared IPMI and network port.

Nutanix Deployment with the Brocade VDX 6740

Depending on the Nutanix model, a single appliance can support up to four nodes. A Nutanix appliance needs three IP addresses per node for the initial installation procedure: one for IPMI, one for the hypervisor (by default, an AHV hypervisor, but it could be ESXi or Hyper-V), and one for the Controller VM (CVM).

In the example deployment shown in Figure 4, Nutanix NX1065- has three nodes.

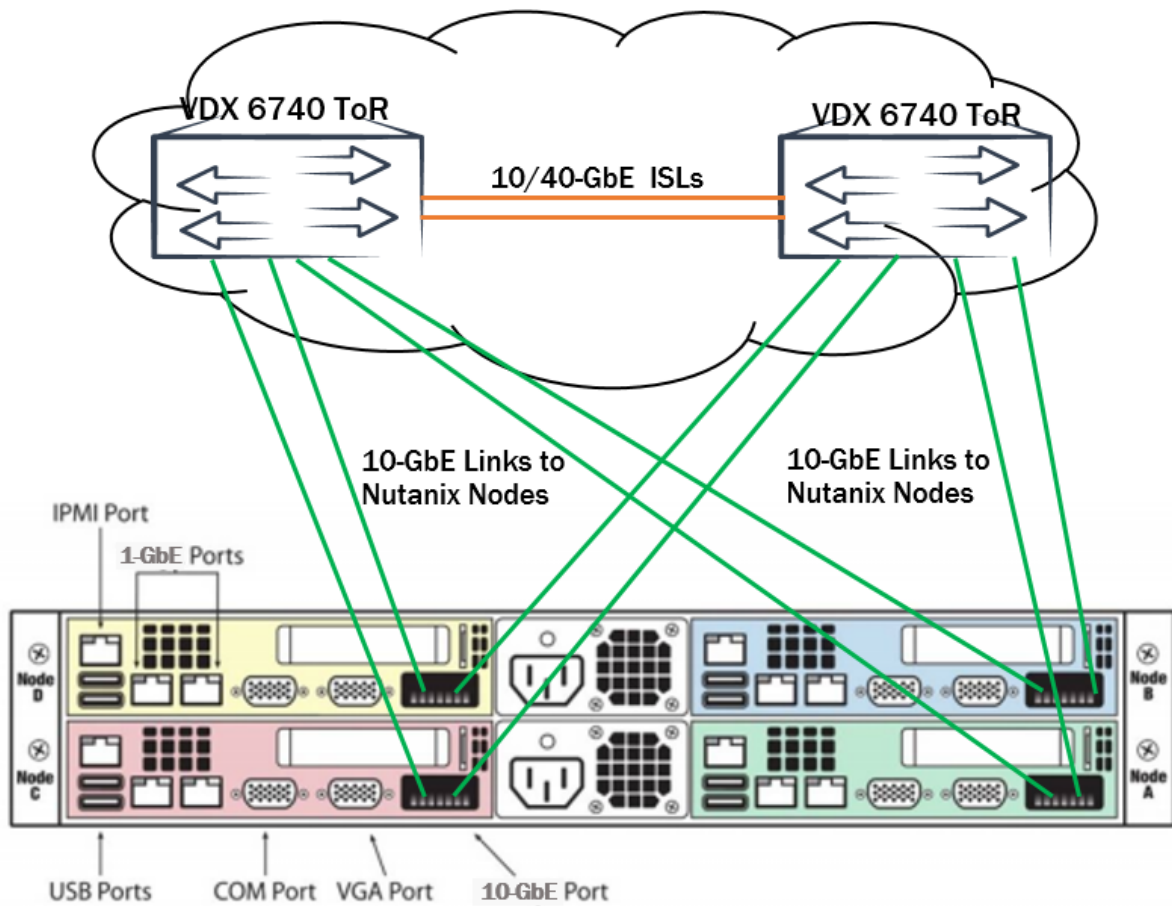
HOST NAME	HOST IP	CVM IP	HYPERVISOR	CPU USAGE	CPU CAPACITY	MEMORY USAGE	MEMORY CAPACITY	TOTAL DISK USAGE
NTNX-155M65230095-A	10.254.4.203	10.254.4.206	KVM	2.04%	28.81 GHz	20.46%	126.13 GiB	3799 GiB of 3.45 TiB
NTNX-155M65230095-B	10.254.4.204	10.254.4.207	KVM	1.84%	28.81 GHz	20.46%	126.13 GiB	3673 GiB of 3.45 TiB
NTNX-155M65230095-C	10.254.4.205	10.254.4.208	KVM	1.74%	28.81 GHz	28.62%	126.13 GiB	34.23 GiB of 3.45 TiB

Figure 4 Example of a 3-Node Nutanix Deployment

Once the imaging process is completed by the Nutanix field engineer, connect the 10-GbE interfaces on the Nutanix nodes according to Figure 5.

Brocade VDX Setup and Configuration

In a typical deployment, the Nutanix nodes are connected to two Brocade VDX 6740 ToR switches in a redundant fashion, as shown in Figure 5. Typically, each Nutanix 2U appliance has 4 nodes, and each node has 2 x 10-GbE connections to the ToR switch.



Nutanix Acropolis (4-Node Appliance)

Figure 5 Nutanix Deployment with Brocade ToR Switches

Brocade recommends configuring the two Brocade VDX switches in Logical Chassis mode for ease of management and troubleshooting. In Logical Chassis mode, a single switch is used to manage the pair of switches. As more switches are added to the fabric, they can also be managed by the single principal switch. However, having the fabric in Logical Chassis mode is not a prerequisite for connecting to Nutanix nodes.

The recommended firmware version is Brocade Network OS 5.0.1a or later. The procedure for upgrading the firmware and configuring the switches in a single fabric in Logical Chassis mode is shown in this section below.

Firmware Upgrade Procedure

Brocade recommends Brocade Network OS version 5.0.1a or later. After removing the switches from their boxes, power each switch on and connect to the console port. If you need to upgrade the switch, you can use the **firmware download usb** command, as shown below. This command loads the firmware from the USB to the system, reboots the system, and commits the firmware automatically.


```
usb on
firmware download usb directory <directory_where_firmware_is_present> coldboot

For example, if you are using a Brocade branded USB, then:

usb on
firmware download usb directory NOS_v5.0.0 coldboot
```

Figure 6 Brocade VDX Firmware Download from USB

Or, if you can use the **firmware download interactive** command to download firmware from an external host or from an attached USB device. You can run this command interactively or provide the parameters on the command line.

Note that if you are using the **firmware download interactive** command, you must set up an in-band or out-of-band management interface to be on the same subnet as your FTP server; make sure that you can reach the FTP server, and then use the following **firmware download** command.

```
sw0# firmware download interactive

firmware download interactive
Server name or IP address: 10.31.2.25

File name: /users/home40/Builds/NOS_v5.0.0

Protocol (ftp, scp, sftp): ftp

User: fvt

Password: *****

Do manual download [y/n]: n

System sanity check passed.
Do you want to continue? [y/n]: y
```

Figure 7 Brocade VDX Firmware Download from FTP Server

Configuring the Two Brocade VDX Switches in a Single Fabric in Logical Chassis Mode

Each switch in an Ethernet fabric is assigned a unique identifier called an RBridge ID. VCS Fabric IDs (also called VCS IDs) identify the fabric membership. Every RBridge that belongs to the same fabric must share the same VCS ID.

In the configuration examples shown below, 10-GbE interfaces are represented as `interface TenGigabitEthernet <rbridge-id/slot/port>`; and 40-GbE interfaces are represented as `interface FortyGigabitEthernet <rbridge-id/slot/port>`. On the Brocade VDX 6740, the slot always remains 0.

Once you remove the switches from the boxes, power them on, and, if needed, upgrade the firmware to Brocade Network OS 5.0.1a or later. Make sure that both switches have the same firmware. Do the following to bring the switches to the Logical Chassis mode.

Enter the following command, where `vcsid` is a unique identifier for the entire fabric, and `rbridge-id` is a unique identifier per switch (by default, the RBridge ID of a Brocade VDX 6740

switch is 1). Make sure to configure unique Rbridges and the same VCS ID since you want to set up the two Brocade VDX switches in a single Logical Chassis mode.

```
sw0# vcs vcsid 1 rbridge-id 1 logical-chassis enable
This operation will perform a VCS cluster mode transition for this local node
with new parameter settings. This will change the configuration to default and
reboot the switch. Do you want to continue? [y/n]:y
```

Figure 8 Brocade VDX—Enabling Logical Chassis Mode on VDX Switch 1

```
sw0# vcs vcsid 1 rbridge-id 2 logical-chassis enable
This operation will perform a VCS cluster mode transition for this local node
with new parameter settings. This will change the configuration to default and
reboot the switch. Do you want to continue? [y/n]:y
```

Figure 9 Brocade VDX—Enabling Logical Chassis Mode on VDX Switch 2

Once the switches are back online after a reboot in Logical Chassis mode, you can enter the **show vcs** and **show fabric isl** commands to make sure that you are in Logical Chassis mode. Configuration output of the VCS cluster is captured below. Note that the VCS cluster is running the Brocade Network OS 5.0.1b GA image. VCS cluster information is also given below. The arrow (>) indicates the cluster principal node. The asterisk (*) indicates the current logged-in mode.

```
sw0# show vcs
Config Mode      : Distributed
VCS Mode        : Logical Chassis
VCS ID          : 10
VCS GUID        : ee691c29-0685-4b82-aa86-9b382722fd04
Total Number of Nodes      : 2
Rbridge-Id      WWN                Management IP   VCS Status
Fabric Status   HostName
-----
1                >10:00:00:27:F8:80:CD:CC*          10.254.1.16     Online
Online          sw0
2                10:00:00:27:F8:81:69:ED          10.254.1.17     Online
Online          sw0

sw0# show fabric isl ports | in Up
64    Te 1/0/1      Up      ISL 10:00:00:27:f8:81:69:ed "sw0"
(downstream)(Trunk Primary)
65    Te 1/0/2      Up      ISL (Trunk port, Primary is 1/0/1 )

sw0# show fabric isl

Rbridge-id: 1    #ISLs: 1

Src      Src      Nbr      Nbr
Index   Interface  Index   Interface      Nbr-WWN      BW
Trunk   Nbr-Name
-----
64      Te 1/0/1      64      Te 2/0/1        10:00:00:27:F8:81:69:ED  20G
Yes     "sw0"
```

```
sw0# show version

Network Operating System Software
Network Operating System Version: 5.0.1
Copyright (c) 1995-2014 Brocade Communications Systems, Inc.
Firmware name:          5.0.1b
Build Time:             20:20:07 Mar  6, 2015
Install Time:           23:27:22 Apr  3, 2015
Kernel:                 2.6.34.6

BootProm:               1.0.1
Control Processor:      e500mc with 4096 MB of memory
```

Slot	Name	Primary/Secondary Versions	Status
SW/0	NOS	5.0.1b 5.0.1b	ACTIVE*
SW/1	NOS	5.0.1b 5.0.1b	STANDBY

Figure 10 Brocade VDX—Show Output Displaying Logical Chassis Mode and Firmware Version

Configuring Brocade VDX Appliance-Facing Ports

The appliance-facing ports on each Brocade VDX switch are connected as trunk ports (with no tagging of frames arriving on native vlan 1). A sample configuration from one of the Brocade VDX switches is shown in Figure 11. Make sure that use this configuration on both Brocade VDX switches.

```
!
interface TenGigabitEthernet 1/0/3
  description Connected-to-Nutanix-Node1
  fabric isl enable
  fabric trunk enable
  switchport
  switchport mode trunk
  switchport trunk allowed vlan all
  no switchport trunk tag native-vlan
  spanning-tree shutdown
  no shutdown
!
interface TenGigabitEthernet 1/0/4
  description Connected-to-Nutanix-Node2
  fabric isl enable
  fabric trunk enable
  switchport
  switchport mode trunk
  switchport trunk allowed vlan all
  no switchport trunk tag native-vlan
  spanning-tree shutdown
  no shutdown
!
interface TenGigabitEthernet 1/0/5
  description Connected-to-Nutanix-Node3
  fabric isl enable
  fabric trunk enable
  switchport
  switchport mode trunk
```

```
switchport trunk allowed vlan all
no switchport trunk tag native-vlan
spanning-tree shutdown
no shutdown
!
```

Figure 11 Brocade VDX 6740 Configuration on Appliance-Facing Ports

Configuring VDX In-band Management

As per Nutanix best practices and typical deployment, the management network is configured in-band (through the Brocade VDX switches) as opposed to a separate network (out-of-band). A sample in-band management configuration is shown in Figure 12.

```
!
interface Vlan 1
  name Management
!
rbridge-id 1
  switch-attributes chassis-name VDX6740
  switch-attributes host-name sw0
  vrf mgmt-vrf
    address-family ipv4 unicast
      ip route 0.0.0.0/0 10.254.0.1
    !
    address-family ipv6 unicast
    !
  !
  interface Ve 1
    vrf forwarding mgmt-vrf
    ip proxy-arp
    ip address 10.254.4.223/20
    no shutdown
  !
!
!
interface TenGigabitEthernet 1/0/15
  description Connected-to-Upstream-network
  fabric isl enable
  fabric trunk enable
  switchport
  switchport mode access
  switchport access vlan 1
  spanning-tree shutdown
  no shutdown
!
```

Figure 12 Brocade VDX 6740 Configuration for In-band Management

Nutanix Setup and Configuration

Once the topology connections and Brocade VDX configuration are in place, and after the Nutanix SE does an initial install of the appliance, you will be given a Cluster Virtual IP address. Log in to the Cluster Virtual IP address of the appliance by entering <https://<Cluster Virtual IP address>:9440>. The login credentials will be entered during the initial install. In this case, it is username “admin” and password “admin”. Upon successful authentication, you will see the Nutanix Prism GUI, as shown in Figure 13 and Figure 14.

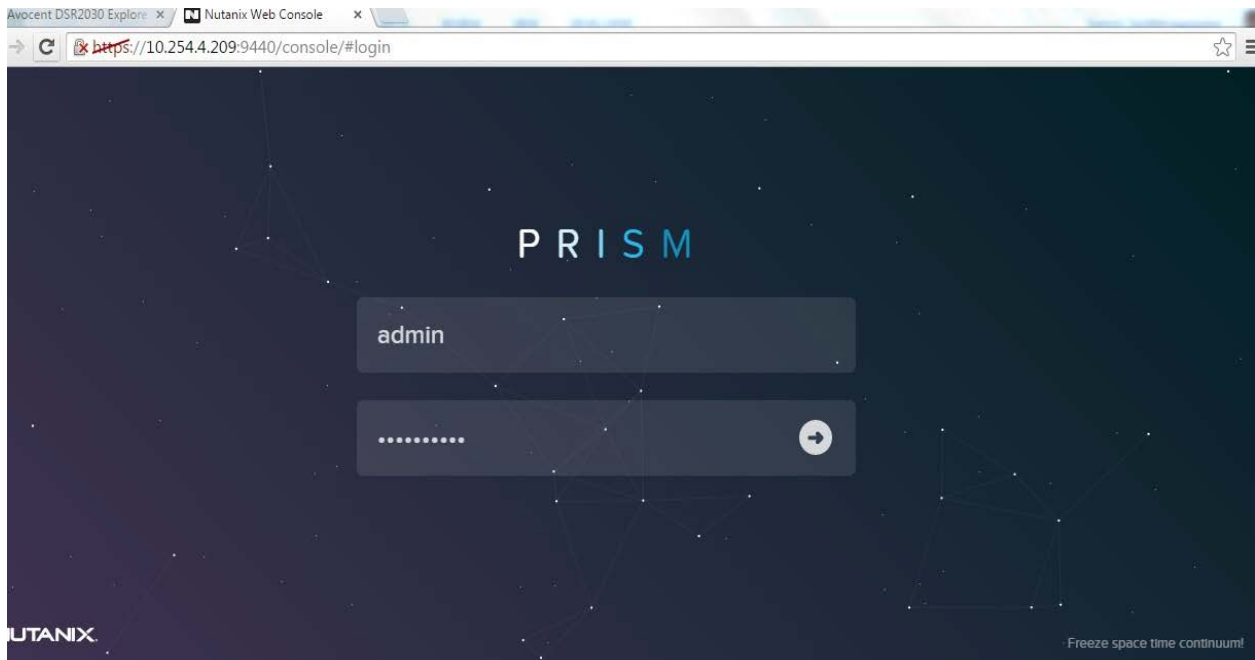


Figure 13 Nutanix Prism GUI Login

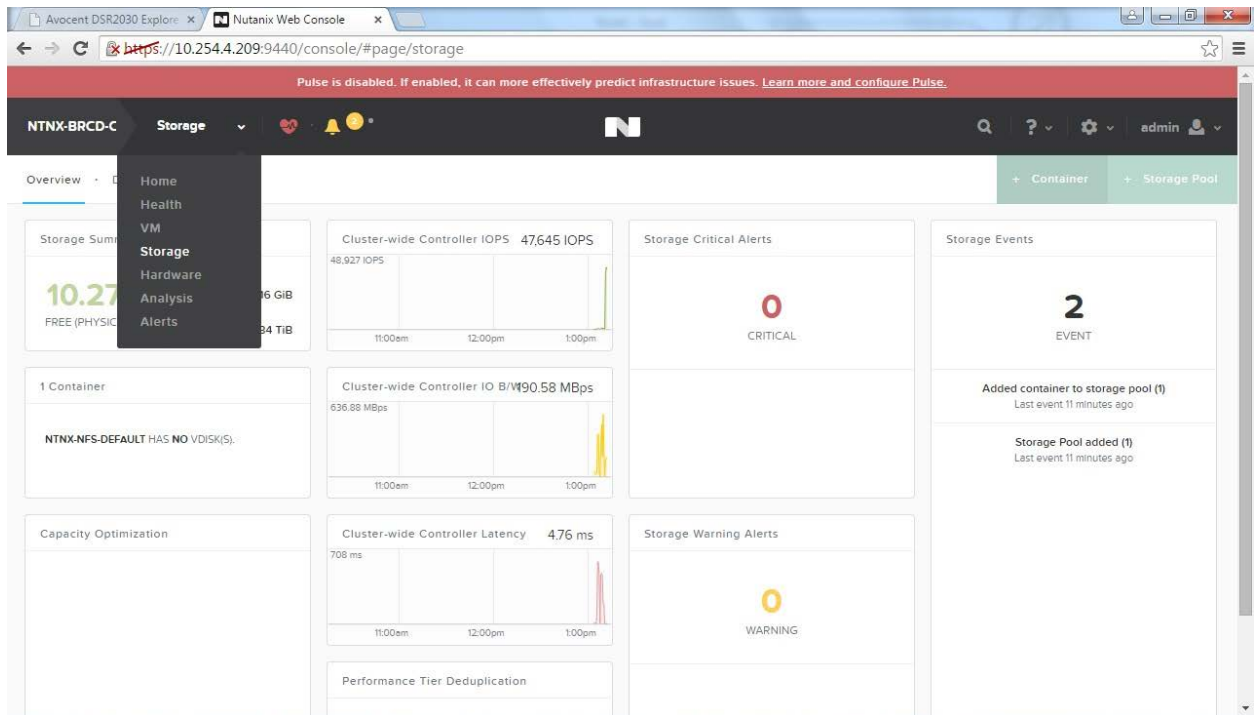


Figure 14 Nutanix Prism GUI Home Screen

Figure 15 and Figure 16 show storage container creation, and Figure 17 shows the network configuration. Once this is complete, you can create application VMs (such as a web server).

Create Container ? X

Enter a name for your container and select a storage pool for it. You can mount the container as an NFS datastore for all hosts, or select individual hosts.

NAME
ISO_files

STORAGE POOL
NTNX-SP-DEFAULT +

MAX CAPACITY
10.25 TiB
(Physical) Based on storage pool free unreserved capacity

NFS DATASTORE
No hosts can be mounted with an NFS datastore.
The cluster's hypervisor(s) are KVM.

⚙️ Advanced Settings Cancel Save

Figure 15 Creating a Container from the Nutanix Prism GUI

NTNX-BRCD-C Storage [Health Icons]

Overview · Diagram · Table

Container Storage Pool 2 Containers

NAME	RF	COMPRESSION	PERF-TIER DEDUPLICATION	ON DISK DEDUPLICATION	ERASURE CODING	FREE (LOGICAL)	USED	MAX CAPACITY
ISO_files	2	Off	Off	Off	Off	5.13 TiB	14.63 GiB	5.14 TiB
NTNX-NFS-DEFAULT	2	Off	None	Off	Off	5.13 TiB	39.84 GiB	5.16 TiB

Figure 16 Default and User-Created Containers in the Prism GUI

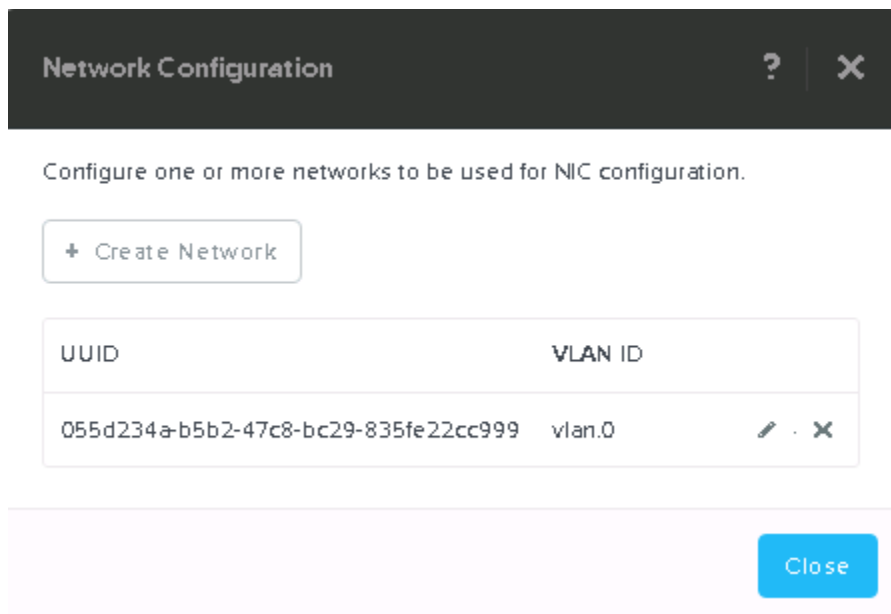


Figure 17 Network Configuration in the Nutanix Prism

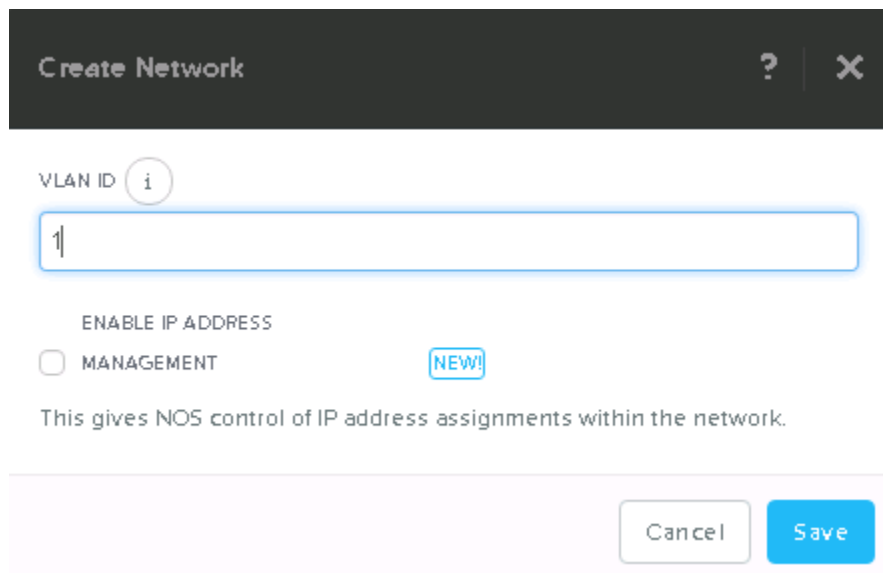


Figure 18 Creating a VLAN ID

Verifying Connectivity Between Nutanix and Brocade VDX

You can verify the neighbor relationship by MAC address lookup. Look for the MAC address of the connected neighbor by using the **show mac-address-table** command on the Brocade VDX, as shown in Figure 19.

```
sw0# show mac-address-table interface ten 2/0/3
Total MAC addresses      : 0
sw0# show mac-address-table interface ten 2/0/4
VlanId  Mac-address      Type      State      Ports
1       0cc4.7a64.1f46    Dynamic  Active     Te 2/0/4
1       5254.0015.65d2    Dynamic  Active     Te 2/0/4
Total MAC addresses      : 2
sw0# show mac-address-table interface ten 2/0/5
VlanId  Mac-address      Type      State      Ports
1       0cc4.7a64.2016    Dynamic  Active     Te 2/0/5
1       5254.001a.5d0a    Dynamic  Active     Te 2/0/5
Total MAC addresses      : 2
```

Figure 19 Show MAC Address Table on the Brocade VDX

On the Nutanix GUI, you can go to the table view of Nutanix nodes and see the details by clicking a node. Click the IPMI IP address (highlighted in Figure 20). This will take you to the IPMI screen, as shown in Figure 21, where you can see the node's MAC address.

Summary > NTNX-15SM65230095-B

HOST DETAILS	
Host Name	NTNX-15SM65230095-B
Hypervisor IP	10.254.4.204
Controller VM IP	10.254.4.207
IPMI IP	10.254.4.201
Node Serial	9e300424-ee15-4bc5-82af-d7f3675bbd7f
Block Serial	15SM65230095
Block Model	NX-1065-G4
Storage Capacity	3.45 TiB
Disks	HDD: 2 disks SSD: 1 disks

Figure 20 Nutanix Node Details from Prism

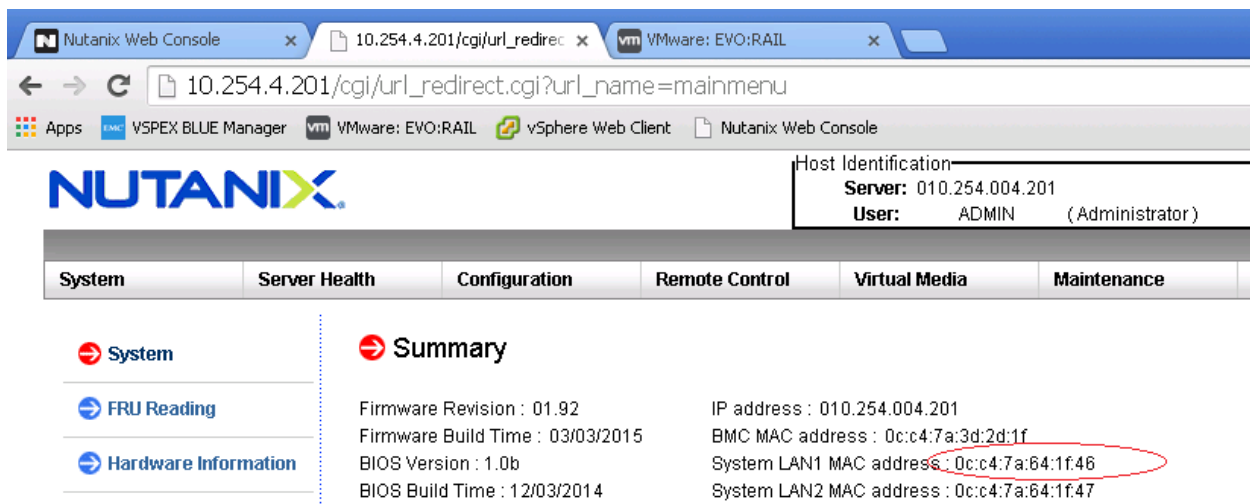


Figure 21 IPMI MAC Address Details from the IPMI Screen

Conclusion

Brocade VCS™ Fabric technology delivers key attributes desired in today's data centers, including simplicity—allowing organizations to manage an entire VCS fabric as a single switch—automated provisioning, multipathing, VM awareness, scale-out, and high performance. It eliminates manual configuration and management, reducing administration and costs. Brocade VCS fabrics enable organizations to scale their networks on demand—without adding complexity—delivering cloud-optimized networking and greater enterprise agility. Due to our native automation, ease of configuration, and operations through our Logical Chassis mode, the Brocade VDX switching platform with our VCS Fabric technology is a perfect complement to a hyperconverged infrastructure.

For more information on Brocade products for IP storage, see:

<http://www.brocade.com/content/brocade/en/products-services/storage-networking/ip-storage.html>

References

1. *Nutanix Acropolis Hypervisor Best Practices*
<http://go.nutanix.com/acropolis-hypervisor-best-practices-guide.html>
2. *Nutanix Tech Note: VMware vSphere Networking on Nutanix*
http://go.nutanix.com/rs/nutanix/images/Nutanix_TechNote-VMware_vSphere_Networking_with_Nutanix.pdf
3. *Nutanix Field Installation Guide*
https://portal.nutanix.com/#/page/docs/details?targetId=Field_Installation_Guide-v3_0:v3_overview_foundation_c.html
4. *Brocade Network OS Administrator's Guide* (for Network OS v6.0.1)
<http://www.brocade.com/content/dam/common/documents/content-types/administration-guide/nos-601-adminguide.pdf>