Dell EMC XC Series Appliances and XC Core Systems

Best Practices for Running VMware ESXi 6.5 or Later Clusters on XC Series Appliances and XC Core Systems

July 2020
Rev. 1.3

DELL EMC

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Contents

1

Revision history

Date	Document revision	Description of changes
July 2020	1.3	Updated the Configuring DRS section.
August 2018	1.2	Updated content to include XC Core.
March 2018	1.1	Updated for second release to add appliances.
November 2017	1	Initial release

2

Introduction

This best practice guidance is aimed at Dell EMC XC Series Appliances and XC Core Systems configured to boot VMware ESXi from either a SATADOM or a Boot Optimized Server Storage (BOSS) card boot device.

This document provides recommendations for maintaining the stability and performance of the platform and workloads, while also preserving the operational lifetime of the boot device.

NOTE: The information in this document applies to both Dell EMC XC Series Appliances, as well as the Dell EMC XC Core System offering. Sections or information that apply to only one of the offerings (XC Series or XC Core) will be called out explicitly.

The Dell EMC XC Series Appliances and XC Core Systems are optimized to host scalable compute, storage, networking, and virtualization workloads. The design focus provides a simplified and scalable approach for handling workloads.

For assistance or questions regarding any of the items that are listed in this document, contact Dell Technical Support.

3

Boot devices

The Dell EMC XC Series Appliances and XC Core Systems support either SATADOM or BOSS (PCIe cards M.2 Drive) as its boot device. The boot device depends on the generation of your appliance.

The 13th generation appliances use SATADOM, and the 14th generation appliances use BOSS. Use the below table to determine which boot device is on your appliance.

Table 1. Associated boot drives per appliance

Appliance	Boot drive
XC430	SATADOM
XC630	SATADOM
XC6320	SATADOM
XC730	SATADOM
XC730xd	SATADOM
XC640	BOSS
XC740xd	BOSS
XC940	BOSS
XC6420	BOSS

Essential information about the boot devices

Dell EMC XC Series Appliances and XC Core Systems differ between 13th and 14th generation appliances.

The 13th generation appliances come with a Serial ATA Disk on Motherboard (SATADOM), which is a flash memory drive designed for use as a boot drive on XC series platforms. While flash memory provides many benefits, it has a finite number of program-erase (P/E) cycles you must consider.

The 14th generation appliances come with a Boot Optimized Server Storage (BOSS) card as the appliance boot device. This PCIe card supports up to two M.2 SATA SSDs configured in RAID1 for high availability.

Both the SATADOM and BOSS cards are designed as appliance boot devices only. Write intensive activities and processes that are leveraged by the Dell EMC XC Series Appliances and XC Core Systems are intended to take place on the SSDs and HDDs, not the boot device itself.

NOTE: The boot device is not intended for application use. Write intensive activities and processes that are leveraged by Dell EMC XC Series Appliances and XC Core Systems are intended to take place on the SSDs and HDDs and not the SATADOM or BOSS boot devices. Any applications defaulting write activity to the BOSS boot drive should be redirected accordingly.

Run virtual machines on the Nutanix Distributed File System only

The boot device is slower performing and more limited in space than the XC Series hosts' SSDs and HDDs used for the highly available Nutanix Distributed File System (NDFS) clustered storage.

Virtual Machines (VMs) run on the boot device are not highly available and potentially fill up the local boot drive, which results in crashing the host hypervisor. This adds additional wear on the boot device.

NOTE: The Nutanix Cluster Checker (NCC) v. 2.2.2 and later will monitor for VMs running on the boot device.

NOTE: A common cause for VMs being run from the boot device is misconfiguration during any add-node or redeployment operation.

When adding or redeploying an ESXi node to the cluster, ensure that the **Virtual Hard Disks (VMDKs)** and **Virtual Machines** locations are configured so that they are directed to the Nutanix Cluster Container location.

4

Configure VMware vSphere vCenter Server

About this task

VMware vSphere vCenter Server should have the following minimum settings defined in a Nutanix environment. In this section, we will outline the steps necessary to properly configure a VMware cluster for a Nutanix environment.

Register vSphere Server under Prism

With the release of AOS 5.0, you can now perform common vCenter Server operations directly from the Prism UI.

About this task

To register vCenter Server with Prism, perform the following operations from the Prism UI:

Steps

1. Navigate to **Settings**.
2. Select **vCenter Registration**.
3. Enter your vCenter Server credentials and click **Register**.

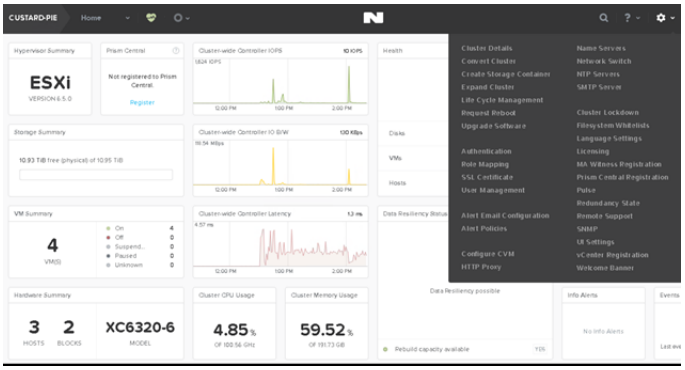

For further details, see the Prism Web Console Guide.

Configure VMware DRS settings

Configure DRS as fully automated, using a conservative policy. The implementation of a more aggressive policy can trigger more frequent vMotion events, which in turn can have an adverse effect on data location.

Prerequisites

1. Disable automation on all CVMs.

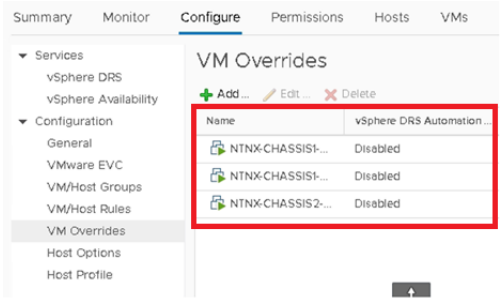

2. Under **vSphere Availability**, disable or turn off **Power Management**.

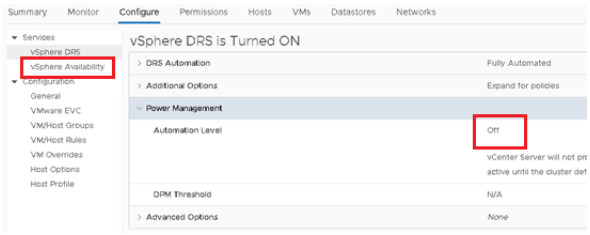

About this task

Use the following steps to configure DRS:

Steps

1. Navigate to **Configure** > **vSphere DRS** and then click **Edit**.
2. Select the **vSphere DRS** check box to enable.
3. From the **DRS Automation** drop-down menu, leave the default migration threshold at 3 and click **OK**.

   This is the default configuration in a fully automated configuration as it is recommended for the Nutanix deployments. This configuration automatically manages data locality so that whenever VMs migrate to another host, writes are always written on one of the replicas locally to maximize the subsequent read performance.

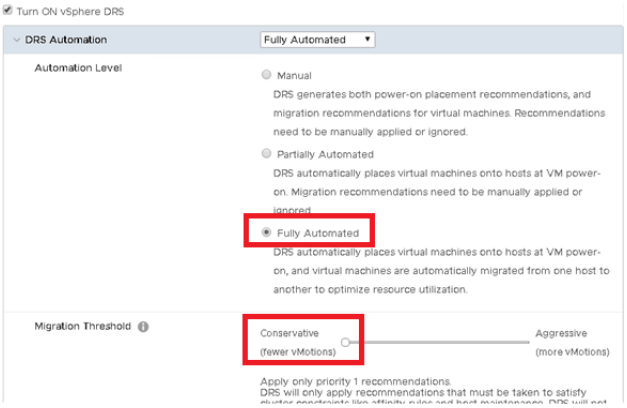

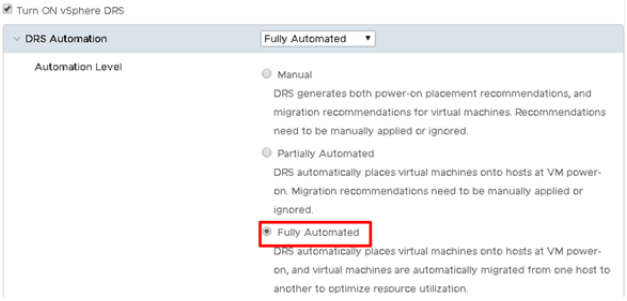

Configure VMware HA

About this task

Use the steps below to configure VMware high availability (HA).

Steps

1. Define the following high availability settings in the cluster:

Setting	Use option
Host Failure Response	Restart VMs
Response for Host Isolation	Power off and restart VMS

Setting	Use option
Admission Control	Refer to vSphere HA Admission Control Settings for Nutanix Environment before enabling this feature
vSphere HA restart Priority for all CVMs	Disabled
VM Monitoring for all CVMs	Disabled
Datastore with PDL	Disabled
Datastore with APD	Disabled

2. Verify that the cluster is now configured as shown below.

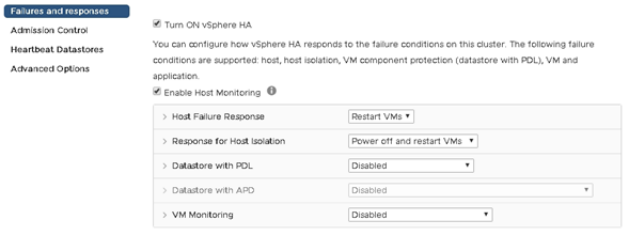

3. **NOTE:**
   - **Known datastore issue**
   - **When configured with a single datastore, an error message in vCenter Server appears stating insufficient datastore defined in the cluster.**
   - **Work around - Add an exception.**

   To address this, add an exception as follows:
   a. Log in to vCenter Server.
   b. Right-click cluster and click **Settings**.
   c. Click **vSphere HA** > **Edit** > **Advanced Options**.
   d. Under **Option**, add an entry for: `das.ignorInsufficientHbDatastore`.
   e. Under **Value**, type **`true`**.

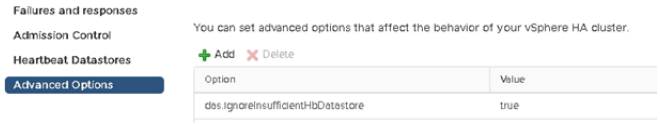

4. Click **Cluster Features**.
5. Clear **Turn on vSphere HA** and click **OK**.
6. Click **Failures & Responses**.
7. Clear **Turn on vSphere HA** and click **OK**.

Enable VMware EVC settings

Enhanced vMotion Compatibility (EVC) simplifies vMotion compatibility issues across CPU generations. Dell EMC recommends enabling EVC so that you can add nodes with newer generation CPUs to existing clusters.

About this task

To enable EVC in the cluster, perform the following steps:

Steps

1. Log in to **vCenter Server**.
2. Right-click the cluster and click **Settings**.
3. Click VMware **EVC** > **EDIT**.
4. Select the appropriate EVC mode based on the CPUs in the existing cluster. Define the rule to set the level to the lowest CPU type in the cluster.

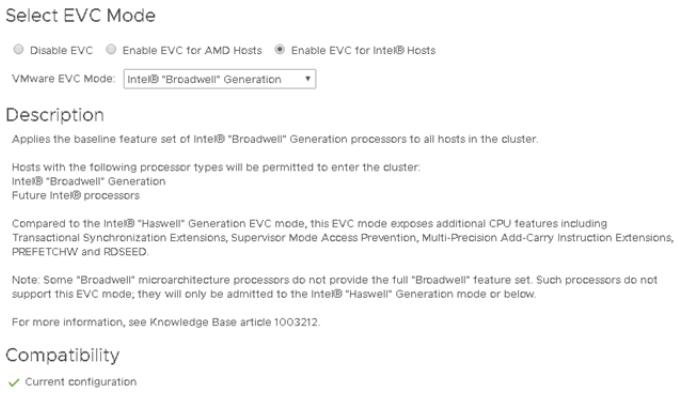

For further details on EVC, see VMware KB 1003212.

5

Apply updates in a VMware vSphere environment

VMware Update Manager

You may use VMware Update Manager (VUM) to apply incremental updates in a Nutanix environment.

Nutanix supports the ability to patch upgrade ESXi hosts with versions that are greater than or released after the Nutanix qualified version, but Nutanix might not have qualified those releases.

When implementing VUM, ensure the following restrictions apply:

- Ensure that you do not import any third-party patches into VUM.
- Ensure that you clear the **Remove 3rd party add-ins** option in the VUM interface.

If other issues occur with the upgrade process, an alert is raised in the **Prism Alert** dashboard. For further information, refer to the vSphere Update manager Installation and Administration Guide for instructions.

Pre-requisite for a One-Click upgrade process

Dell EMC recommends that ESXi host updates are applied through Prism's One-Click feature, as this is a cluster-aware process.

About this task

Nutanix qualifies specific VMware ESXi hypervisor updates and provides a related JSON metadata upgrade file on the Nutanix Support Portal for one-click upgrade through the Prism web console Software Upgrade feature.

Dell EMC provides qualified ESXi binary files to Nutanix, and Nutanix in turn provide the related JSON metadata upgrade files.

Steps

1. Log in to the Nutanix portal with your credentials.
2. Go to **Downloads** > **Hypervisor Details** and download the JSON file for the ESXi release that you want to apply to the cluster nodes.
3. Download the appropriate ESXi binary .zip file from the Dell support site.
4. Run the Nutanix Cluster Checks (NCC) using one of the following methods:
   - From the Prism web console **Health** page, select **Actions** > **Run Checks** > **All Checks** and then click **Run**.
   - Log in to a Controller VM and use the `ncc CLI`

     ```
     nutanix@cvm$ ncc health_checks run_all
     ```

     NOTE: If the check reports a status other than `PASS`, resolve the reported issues before proceeding. If you are unable to resolve the issues, contact Dell Support for assistance.
5. Disable Admission Control within vCenter Server, on performing the following:
   a. Right-Click on the cluster resource and select **Settings**.
   b. Select **Services** > **vSphere Availability** > **EDIT**.
   c. Select **Admission Control** > **Define host failover capacity by** and set it to **Disabled**.

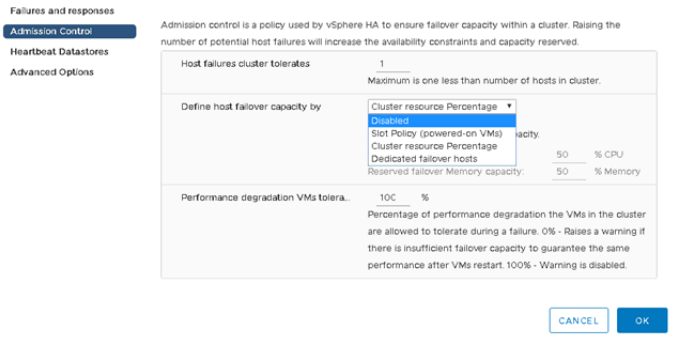

Perform the upgrade

Steps

1. Within Prism, click the gear icon in the upper right corner of the interface and select **Upgrade Software**.

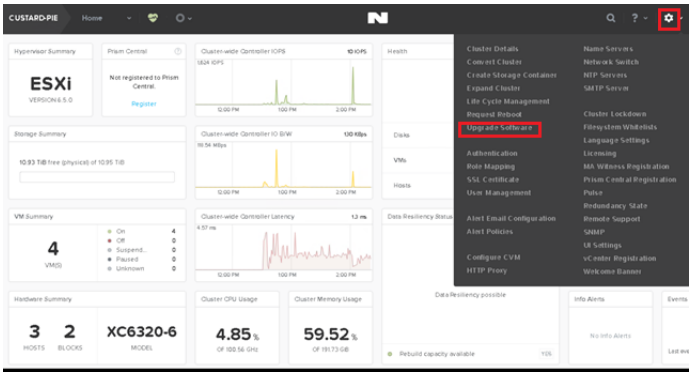

2. On the **Hypervisor** tab, upload the JSON file, and the ESXi bundle and then click **Upload Now**.

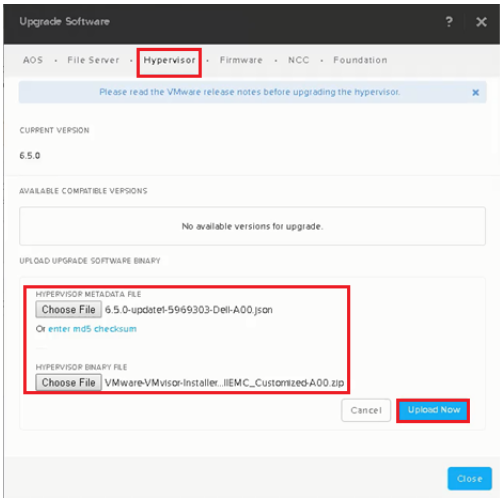

3. **Pre-upgrade** is an optional process to validate the configuration on the CVM on which you are connected to before proceeding. These checks also run as part of the upgrade procedure.

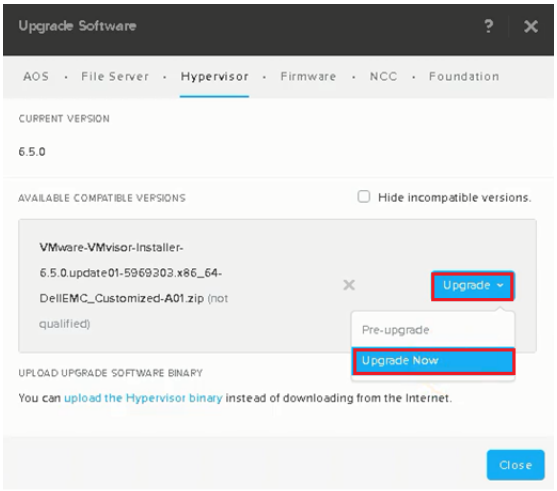

4. When the upload process completes, click **Upgrade** > **Upgrade Now**, then click **Yes** to confirm.
5. Enter in the VMware vCenter Server credentials when prompted and click **Upgrade**.

   The Upgrade Software dialog box shows the progress of your selection, including status of pre-installation checks and uploads, through the Progress Monitor. All VMs running on the node will be migrated to another node in the cluster before the node is placed into maintenance mode. Once the node has completed the upgrade process, it will be re-added to the cluster before the process moves on to the remaining nodes.
6. When all nodes have been successfully upgraded, repeat step 4 to validate that no errors have occurred during the process.
7. Re-enable **Admission Control** within vCenter Server with the following:
   a. Right-Click on the cluster resource and select **Settings**.
   b. Go to **Services** > **vSphere Availability** > **EDIT**.
   c.
   d. Re-enable your previously defined settings.

6

Deployment best practices

Virtual disk provisioning

Dell EMC recommends using thin provisioning disks in an XC Series environment, as thick eager-zero virtual disks offer no performance benefits over a thin virtual disk.

For further information, see Nutanix KB 1591 (Log in required).

SYSLOG

Except for VMware NSX implementations, the configuration of a remote SYSLOG server is not an absolute requirement; however, there are benefits and use cases for such an implementation. For example:

- Central monitoring of ESXi hosts.
- Persistent logging in situations where the boot device becomes inaccessible or read-only.

See VMware KB 2003322 for detailed information on SYSLOG configuration.

NOTE: Dell EMC does not recommend moving the ESXi scratch partition to Nutanix NFS, as this may become inaccessible in the event of a cluster stop event.

A

Appendix

Additional resources and references

Support.dell.com is focused on meeting your needs with proven services and support.

Referenced or recommended Dell EMC publications

For Dell EMC XC Series Appliances and XC Core Systems product documentation, go to *Dell.com/xcseriesmanuals*.

Referenced or recommended Nutanix publications

- Nutanix Cluster Check (NCC) 2.2.x Guide (Log in required)
- Nutanix KB 2490 - NCC Health Check: host_disk_usage_check (Log in required)
- Nutanix KB 3196 - NCC Health Check: sata_dom_uvm_check (Log in required)
- Create vmware Nutanix internal vswitch from Scratch in ESXi (Log in required)
- Perform Updates with Life Cycle manager (Log in required)
- VMware vSphere Networking on Nutanix
- Prism Web Console Guide
- The Nutanix Bible