

Surveillance Dell EMC Storage with Qognify

Configuration Guide

H14189

REV 1.3



Copyright © 2015-2017 Dell Inc. or its subsidiaries. All rights reserved.

Published November 2017

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS-IS." DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.
Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Chapter 1	Introduction	5
	Purpose.....	6
	Scope.....	6
	Assumptions.....	6
Chapter 2	Configuring the solution	9
	Design concepts.....	10
	EMC VNX.....	11
	Disk drives.....	11
	Storage pool configuration (recommended).....	11
	LUN configuration.....	12
	Fibre Channel configuration.....	13
	Dell EMC PowerPath.....	13
	FC LUN configuration for a virtualized environment.....	14
	iSCSI initiators.....	14
	Configure iSCSI front-end ports	14
	Connect the iSCSI target on Windows.....	15
	Format the iSCSI target on Windows.....	15
	Recommended cache configuration.....	16
	Isilon (NAS).....	16
	Volume limits.....	16
	Large file system, small view (SmartQuotas).....	16
	Configuring SmartQuotas (recommended).....	17
	Unique share naming.....	17
	Configuring SmartConnect (optional).....	18
	SMB specific configuration.....	19
	Link aggregation.....	21
	I/O optimization configuration.....	22
	Configuring authentication and access control.....	22
	Releases tested.....	23
	Continuous Availability.....	23
	Job Engines and performance impact.....	24
	SSD strategies.....	25
	DNS specific configuration.....	25
	Manually re-balancing recorders across nodes.....	26
	Network adapter configuration.....	26
	Add a NAS drive	27
	Set the Network storage option	27
	Snapshots and data progression.....	28
	Microsoft Multipath I/O.....	28
	Spare disks.....	28
Chapter 3	Conclusion	29
	Summary.....	30
	EMC VNX arrays.....	30
	Dell EMC Isilon scale-out storage.....	30

CONTENTS

CHAPTER 1

Introduction

This chapter presents the following topics:

- [Purpose](#).....6
- [Scope](#).....6
- [Assumptions](#).....6

Purpose

This configuration guide aims to help Dell EMC field personnel understand how to configure Dell EMC storage system offerings to simplify the implementation of Qognify . This document is not a replacement for the Qognify implementation guide nor is it a replacement for the *Dell EMC Storage with Qognify : Sizing Guide*.

Qognify VisionHub 4.2 and NiceVision Net 3.1 provides video surveillance solutions with video enhancement tools, integrated video analytics, and an open, IT-friendly design that delivers extreme performance. With Dell EMC storage systems, Qognify offers reliable smart IP video surveillance packages that are designed to work in and for every environment.

Scope

This guide is intended for internal Dell EMC personnel and qualified Dell EMC and Qognify partners. It provides configuration instructions for installing the Qognify video management software using Dell EMC storage platforms.

The following Dell EMC storage systems have been tested:

- Dell EMC Isilon™
- EMC VNX™

This guide supplements the standard *EMC VNX Storage Best Practices with Video Management Systems: Configuration Guide* and *Dell EMC Isilon Storage Best Practices with Video Management Systems: Configuration Guide* and provides configuration information specific to Qognify .

Note

All performance data in this guide was obtained in a rigorously controlled environment. Performance varies depending on the specific hardware and software used.

Assumptions

This solution assumes that internal Dell EMC personnel and qualified Dell EMC partners are using this guide with an established architecture.

This guide assumes that the Dell EMC partners who intend to deploy this solution are:

- Associated with product implementation
- Qognify-certified to install Qognify services
- Proficient in installing and configuring VNX storage solutions
- Proficient in installing and configuring Isilon storage solutions
- Familiar with installing and configuring VMware hypervisors and the appropriate operating system, such as Microsoft Windows or a Linux distribution
- Able to access the *EMC VNX Storage with Video Management Systems: Configuration Guide* and *Dell EMC Isilon Storage with Video Management Systems: Configuration Guide*

The configurations that are documented in this guide are based on tests that we conducted in the Dell EMC Surveillance Lab using worst-case scenarios to establish a

performance baseline. Lab results might differ from individual production implementations.

CHAPTER 2

Configuring the solution

This chapter presents the following topics:

• Design concepts	10
• EMC VNX	11
• Isilon (NAS)	16
• Releases tested	23
• Continuous Availability	23
• Job Engines and performance impact	24
• SSD strategies	25
• DNS specific configuration	25
• Manually re-balancing recorders across nodes	26
• Network adapter configuration	26
• Add a NAS drive	27
• Set the Network storage option	27
• Snapshots and data progression	28
• Microsoft Multipath I/O	28
• Spare disks	28

Design concepts

There are many design options for a Qognify implementation. Qognify offers many documents and materials related to design and implementation of Qognify . These design details are beyond the scope of this paper.

The software solution is composed of video recorders, a server and client solution, and video analytics.

There are many design options for a NiceVision Net implementation. Qognify offers many documents and materials related to design and implementation of Qognify NiceVision Net. These design details are beyond the scope of this paper.

The NiceVision Net software solution is composed of smart video recorders, a Qognify ControlCenter solution, and video analytics. The solution also includes an extensive software development kit, which enables the integration of third-party security systems with NiceVision Net. NiceVision Net 3.1 provides a smooth migration path from analog to IP/digital technologies, with a video management offering that improves the performance of the analog cameras.

NiceVision Smart Video Recorders (SVR) are complete, high-performance network recording, and video management solutions. The video recording, video value-added services, and analytics can turn any channel into a smart one. The SVR family is fully scalable, can manage encoders with third-party IP devices, and offers a migration path from analog to IP.

NiceVision encoders enable you to seamlessly record, manage, and analyze highquality 4CIF real time video images from analog cameras over an IP network. NiceVision Video Analytics provides proactive alerts to potential unfolding events using applications for Perimeter Intrusion Detection, Crowd Management, and Situation Indication, for improved response.

The components of the NiceVision ControlCenter management are:

Virtual Matrix (VMX)

Cost-effective alternative to an analog matrix.

Event management

Real time monitoring and advanced investigation.

Web and smartphone solution


Enables on-the-move security on any web-based device.

Web deployment

For easy installation and launch of ControlCenter applications.

Tests were conducted using physical servers for Qognify SVR. In the Dell EMC Surveillance Lab environment, five node Isilon clusters were used for testing.

The following figure represents the basic configuration that was tested in our lab.

Figure 1 Qognify architecture

EMC VNX

VNX storage is ideal for recording and managing terabytes of video from distributed locations. This section describes best practices for configuring a VNX storage system for this solution.

The VNX family includes the VNX and VNX-VSS series arrays. The VNX series is designed for midtier to enterprise storage environments, is ideal for distributed environments, and can scale to handle large petabyte (PB) environments with block-only requirements at central locations.

Disk drives

Although any supported drive will work, video surveillance systems typically rely on the density of the array. Dell EMC recommends NL-SAS drives of the highest available density in this solution. In general, we used one-terabyte (TB) or multi-TB NL-SAS drives when performing our tests.

Note

Because of the high percentage of sequential, large block writes, Dell EMC does not recommend using flash drives for video storage within a surveillance application.

Storage pool configuration (recommended)

The tests we conducted show how storage pools that are defined with the maximum allowable number of disks per pool perform as well as, or better than, traditional RAID groups. Therefore, Dell EMC recommends that you use storage pools rather than RAID groups. Storage pools also reduce the required array management tasks.

The VNX family array architecture is optimized for storage pools. A storage pool is a construct that is built over one, or more commonly multiple, RAID groups. LUNs are built on top of the storage pool. The read/write activity is a random distribution across all disks defined to the storage pool. This distribution results in increased and balanced per disk utilization and improved performance when compared to traditional RAID implementations.

The RAID groups underlying storage pools can be either RAID 5 or RAID 6. The default and recommended RAID configuration for a VNXe or VSS1600 array using NL-SAS drives is RAID 6. Either RAID 5 or RAID 6 can be used with VNX arrays. RAID 5 is used for optimizing the array to achieve the maximum amount of storage and RAID 6 is

used for enhancing data protection. Our tests using an isolated surveillance infrastructure did not reveal any notable performance variances when using RAID 5 as compared to RAID 6.

Building a storage pool is a straightforward process. You can configure either RAID 5 or RAID 6 pools depending on the VNX storage system restrictions and the level of risk that the customer is willing to accept. When configuring storage pools, use large storage pools with large logical unit number (LUN) sizes, and configure the LUNs as thick. Do not use thin LUN provisioning.

Dell EMC recommends the following RAID configurations for VNX arrays:

- RAID 5 or RAID 10 with SAS drives
- RAID 6 with NL-SAS drives

Procedure

1. In Unisphere, select **Storage > Storage Pools** for block.
2. Click **Create** under **Pools** in the **Pools** section.
3. Set the following options for the storage pool:
 - Storage pool name
 - RAID type
 - Number of SAS drives
 - Number of NL SAS drives
4. Choose a method for selecting disks to include in the storage pool:
 - **Automatic:** Provides a list of available disks.
 - **Manual:** Enables you to select specific disks to include in the storage pool from a list of available disks. Be sure to clear the automatic disk recommendation list before you select new disks from the list.
5. Select **Perform a Background verify on the new storage** and set the priority to medium.
6. Click **Apply**, and then click **YES** to create the storage pool.

LUN configuration

A VNX pool LUN is similar to a classic LUN. Pool LUNs comprise a collection of slices. A slice is a unit of capacity that is allocated from the private RAID groups to the pool LUN when it needs additional storage. Pool LUNs can be thin or thick.

Thin LUNs typically have lower performance than thick LUNs because of the indirect addressing. The mapping overhead for a thick LUN is less than for a thin LUN.

Thick LUNs have more predictable performance than thin LUNs because they assign slice allocation at creation. Because thick LUNs do not provide the flexibility of oversubscribing like a thin LUN, use thick LUNs for applications where performance is more important than saving space.

Thick and thin LUNs can share the same pool, enabling them to have the same ease-of-use and benefits of pool-based provisioning.

Procedure

1. In Unisphere, right-click a storage pool and then click **Create LUN**.
2. Type the user capacity for the LUN.

3. Type the starting **LUN ID**, and then select the number of LUNs to create.

For example, if the selected LUN ID is 50, and the selected number of LUNs to create is 3, the names for the LUNs are 50, 51, and 52.

4. Select **Automatically assign LUN IDs as LUN names**.
5. Click **Apply**.

Fibre Channel configuration

To transfer traffic from the host servers to shared storage, the serial-attached network (SAN) uses the Fibre Channel (FC) protocol that packages SCSI commands into FC frames.

Note

iSCSI is prevalent for video security implementations because it often provides a lower-cost option when compared to FC.

To restrict server access to storage arrays that are not allocated to the server, the SAN uses zoning. Typically, zones are created for each group of servers that access a shared group of storage devices and LUNs. A zone defines which HBAs can connect to specific service providers (SPs). Devices outside a zone are not visible to the devices inside the zone.

Zoning is similar to LUN masking, which is commonly used for permission management. LUN masking is a process that makes a LUN available to some hosts and unavailable to other hosts.

Zoning provides access control in the SAN topology. Zoning defines which HBAs can connect to specific targets. When you use zoning to configure a SAN, the devices outside a zone are not visible to the devices inside the zone.

Zoning has the following effects:

- Reduces the number of targets and LUNs presented to a host
- Controls and isolates paths in a fabric
- Prevents non-ESXi systems from accessing a particular storage system and from possible virtual machine file system (VMFS) data loss
- Optionally, separates different environments, such as test and production environments

With VMware ESXi hosts, use single-initiator zoning or single-initiator-single-target zoning. The latter is the preferred zoning practice because it is more restrictive and prevents problems and misconfigurations that can occur on the SAN.

Dell EMC PowerPath

Dell EMC PowerPath™ is recommended for block storage (FC and iSCSI) implementations. PowerPath Multipathing automates data path management, failover and recovery, and optimizes load balancing to ensure application availability and performance.

FC LUN configuration for a virtualized environment

Fibre Channel LUNs can be configured using two methods for virtualized environments.

Raw device mappings (RDM)

RDM can be used for virtual machine archivers to store Video data. The RDM allows a virtual machine to access and use the storage device directly. The RDM contains metadata for managing and redirecting disk access to the physical device. The file gives you some of the advantages of direct access to a physical device while keeping some advantages of a virtual disk in VMFS. As a result, it merges VMFS manageability with raw device access.

FC datastores

ESXi uses datastores, which are logical containers that hide specifics of physical storage from virtual machines and provide a uniform model for storing virtual machine files. Datastores that you deploy on block storage devices use the vSphere VMFS format, a special high-performance file system format that is optimized for storing virtual machines.

iSCSI initiators

Software or hardware initiators may be used with VMware ESXi server or a non-virtualized server.

Microsoft Internet SCSI (iSCSI) initiators

For both physical servers and VMware ESXi server, the Dell EMC Surveillance Lab uses Microsoft iSCSI initiators with excellent results.

Hardware iSCSI initiators

Hardware iSCSI initiators can be used. There are many iSCSI initiators available on the market, and results might vary.

Configure iSCSI front-end ports

Configure the storage system iSCSI front-end ports when the cabling is completed.

For cable specifications, refer to the technical specifications for your storage system. You can generate an up-to-date version of these specifications using the **Learn about storage system** link on the storage system support website.

For high availability:

- Connect one or more iSCSI front-end data ports on SP A to ports on the switch or router. If two switches or routers are available, connect the same number of iSCSI front-end data ports on SP B to ports on the same switch or router, or on another switch or router.
- For a multiple NIC or iSCSI host bus adapter (HBA) server, connect one or more NIC or iSCSI ports to ports on the switch or router. If two switches or routers are available, connect the same number of NIC or iSCSI HBA ports to ports on the same switch or router, or on another switch or router.

Procedure

1. To start Unisphere, in an Internet browser, type the IP address of the storage system SP that you want to manage.

2. Type your user name and password.
3. Click **Login**.
4. From Unisphere, select **System > Hardware > Storage Hardware**.
5. Identify the storage system iSCSI front-end ports by selecting **SPs > SP A/B > IO Modules > Slot [#] > Port [#]** in the **Hardware** window.
For example: **SPs > SP A > IO Modules > Slot A4 > Port 0**
6. Click **OK**.
7. Highlight the iSCSI front-end port that you want to configure and click **Properties**.
8. To assign an IP address to the port, click **Add** in **Virtual Port Properties**.
9. Click **OK** and close all open dialog boxes
10. Click **OK**.
11. Click **OK**.

The iSCSI Port Properties window displays the added virtual ports in the Virtual Port Properties area.

Connect the iSCSI target on Windows

When the iSCSI target is connected to the Windows iSCSI initiator, the volume is shown on the computer as a local physical hard drive, which can be used for video storage.

Procedure

1. Connect the iSCSI target with the Windows iSCSI initiator.
 - a. Launch the iSCSI initiator at **Control Panel > Tools**.
 - b. On the iSCSI Initiator **Properties** page, click **Discovery**.
 - c. Enter the IP address of the NAS and then click **OK**.
 - d. Click **Targets** and then select the available iSCSI targets that you want to connect.
 - e. Click **Connect**.
 - f. Click **OK**.

On successful connection, the status changes to Connected.

Format the iSCSI target on Windows

After the iSCSI target has been successfully connected on Windows, Windows displays the iSCSI target as an Unallocated Disk. You must set the disk to online and format the disk before you can start using it as a local disk to store video.

Procedure

1. Right-click **Computer** and then click **Manage**.
2. Click **Disk Management** to display current disk information.
3. Right-click **iSCSI Disk** and then click **Online** to activate the disk.
4. Right-click **iSCSI Disk** again to open the **New Simple Volume Wizard** window.
5. Follow the wizard to complete formatting the disk.

When the wizard completes, the disk appears as a local hard disk drive, which can then be used as extra storage space.

Recommended cache configuration

EMC VNX generation 2 systems, such as VNX5200 or VNX5400, manage the cache. If the array is shared with other applications, you can use a lower write cache value, but avoid excessive forced flushes.

Dell EMC recommends that you configure the cache as 90 percent write and 10 percent read if the storage array does not automatically adapt to the write characteristics of video surveillance (for example, EMC VNX5500 or EMC VNX-VSS100).

Isilon (NAS)

The Isilon scale-out network-attached storage (NAS) platform combines modular hardware with unified software to harness unstructured data. Powered by the distributed Isilon OneFS™ operating system, an Isilon cluster delivers a scalable pool of storage with a global namespace.

The platform's unified software provides centralized web-based and command-line administration to manage the following features:

- A symmetrical cluster that runs a distributed file system
- Scale-out nodes that add capacity and performance
- Storage options that manage files and tiering
- Flexible data protection and high availability
- Software modules that control costs and optimize resources

Volume limits

Implementations greater than 8 TB are common when video is stored on high-end storage, such as Isilon scale-out NAS storage and VNX block storage. The clustered file system OneFS uses enables Isilon to handle these large volumes.

Large file system, small view (SmartQuotas)

Although it is possible to assign the full Isilon cluster file system to a single Qognify SVR, the Dell EMC best practice is to use SmartQuotas™ to segment the single Isilon file system so that each SVR has a logical subset view of storage.

There are three directory-level quota systems:

Advisory limit

Lets you define a usage limit and configure notifications without subjecting users to strict enforcement.

Soft limit

Lets you define a usage limit, configure notifications, and specify a grace period before subjecting users to strict enforcement.

Hard limit (recommended)

Lets you define a usage limit for strict enforcement and configure notifications. For directory quotas, you can configure storage users' view of space availability as reported through the operating system.

Use the **Hard limit** quota system to set the video storage as a defined value.

If necessary, both Isilon and the Qognify SVR can add or subtract storage, even if a hard quota is set.

Configuring SmartQuotas (recommended)

The SmartQuotas feature enables you to limit the storage that is used for each Qognify SVR. It presents a view of available storage that is based on the assigned quota to the SVR. SmartQuotas enables each SVR to calculate its available disk space and react appropriately.

Without SmartQuotas, the Qognify administrator must anticipate the total write rate to the cluster and adjust the **Min Free Space** on each SVR accordingly. A miscalculation can result in lost video. SmartQuotas resolves the issues that can be caused by manual calculations.

Configure SmartQuotas when more than one SVR is writing to the Isilon cluster, or when other users share the cluster. Enable SmartQuotas and define a quota for each share or directory.

Configure the SmartQuotas setup with the following settings:

- Configure a hard share limit threshold to the SVR video files.
- Define OneFS to show and report the available space as the size of the hard threshold.
- Set the usage calculation method to show the user data only.

Procedure

1. From the OneFS GUI, select **File System Management > SmartQuotas**.
2. For each listed share, select **View details**.
3. Under **Usage Limits**, select **Edit usage limits**.
4. Define the SmartQuotas limit and set the threshold:
 - a. Select **Specify Usage Limits**.
 - b. Select **Set a hard limit**.
 - c. Type the hard limit value.
 - d. Select the size qualifier, typically **TB**.
 - e. Select the size of the hard threshold.
5. Click **Save**.
6. Repeat the process for the remaining shares.

Unique share naming

When working with a single file system, each SVR uses the time and date as part of its directory and file-naming conventions.

To avoid corruption caused by overwriting or grooming (deleting) files prematurely, create a unique share for each SVR.

Configuring SmartConnect (optional)

SmartConnect™ uses the existing Domain Name Service (DNS) Server and provides a layer of intelligence within the OneFS software application.

The resident DNS server forwards the lookup request for the delegated zone to the delegated zone's server of authority, which is the SmartConnect Service IP (SIP) address on the cluster. If the node providing the SmartConnect service becomes unavailable, the SIP address automatically moves to a different node in the pool.

Connections are balanced across the cluster, which ensures optimal resource utilization and performance. If a node goes down, SmartConnect automatically removes the node's IP address from the available list of nodes, ensuring that a connection is not tried with the unavailable node. When the node returns to service, its IP address is added to the list of available nodes.

The delegated server authority is always the node with the lowest ID, unless it has surrendered its authority status, either voluntarily or involuntarily. This node should always be available, but if the status of the node changes and becomes unavailable, it voluntarily surrenders its role as server of authority.

You must add a delegation Name Server (NS) entry to the resident DNS server for the SmartConnect name, which points to the SIP address as the Name Server. In your DNS Manager, create a **New Delegation** using your SmartConnect zone name. In the Microsoft DNS wizard, a New Delegation record is added in the forward lookup zone for the parent domain.

SmartConnect balances connection loads to the Isilon cluster and handles connection failover. With SmartConnect, all Qognify SVRs use a single fully qualified domain name (FQDN) or universal naming convention (UNC) path for video storage access. Using this network name provides load balancing when the connection to the cluster is made and simplifies installations.

SmartConnect Basic can use a round-robin-type connection allocation, which is based on DNS load balancing.

SmartConnect Advanced can include multiple pools for each subnet, Dynamic IP addresses for NFS, and the following load-balancing options (Connection policy and Rebalance policy):

Round-robin (recommended)

Sequentially directs a connection to the next Isilon IP address in the cycle. Based on field reports, this option works well with 20 servers or more.

Connection count

Provides uniform distribution of the Qognify SVR servers to specified nodes in the Isilon cluster. Use a unique IP address pool for video recording and SVR read/write access.

Network throughput

Based on NIC utilization. Use of throughput requires that each SVR is activated, configured, and recording video after it connects to Isilon.

CPU usage

Uses the node CPU utilization to determine which Isilon IP address to assign to the next connection request.

Ensure that no other service uses the SVR IP address pool. Define additional pools for management (such as Isilon InsightIQ™ or administrative access), evidence repository, post process, or other use.

Procedure

1. Select **Networking Configuration**.
2. Under **Subnet > Settings**, define the SmartConnect service IP (SSIP) address. The SSIP address is the IP address that the DNS uses for the Isilon Authoritative name service.
3. Under **Pool settings**:
 - a. Define the SmartConnect zone name, which is the name to which clients connect.
 - b. Define the SmartConnect service subnet (the subnet that has the SSIP configured on the DNS server).
 - c. Define the connection balancing policy to **Round Robin**.
 - d. Set the IP allocation strategy to **Static**.
4. Verify this configuration on the SmartConnect dashboard.

SMB specific configuration

During testing in the Dell EMC Surveillance Lab, we encountered a network connectivity failure issue between the Isilon and video server that lead to a `File Open` issue. The TCP socket connections that were previously made between the video server and the Isilon node were not closed. As a result, the Qognify SVR failed to write to the Isilon share as the files were being opened, and were then not available for further modifications. When SmartConnect was setup and in place, the expected behavior, if the failure is on the Isilon end, was that the connection would move to the next available node.

We worked with the Isilon support team to discover that the TCP socket connections were causing the recovery issue from a network connectivity failure. In the Dell EMC Surveillance Lab, we tested the workaround to keep the socket connection open for a minimum of one minute only, and then closed the socket if the previously connected IP address was not available. This workaround was implemented by adding two timeouts, `keepidle` and `keepintvl`, on the Isilon cluster. The Isilon Development and Support team recommend that we set `keepidle` to 61 seconds, with one minute being the minimum we can assign to this parameter, and `keepintvl` to 5 seconds. Using this configuration, the Qognify SVRs start writing to the share with a data loss interval of 1-2 minutes.

To make a `sysctl` configuration change persistent, add to or change the desired parameter in the `sysctl.conf` file.

Procedure

1. Open an SSH connection on a node in the cluster and log on using the `root` account.
2. Run the following command to back up the `/etc/mcp/override/sysctl.conf` file:

```
touch /etc/mcp/override/sysctl.conf && cp /etc/mcp/override/
sysctl.conf /etc/mcp/override/sysctl.conf.bkul
```

3. Run the following command, where `<sysctl_name>` is the parameter you want to add or change and `<value>` is the value assigned to the parameter.

```
isi_sysctl_cluster <sysctl_name>=<value>
```

The following output is displayed:

```
Value set successfully
```

For example:

```
isi_sysctl_cluster net.inet.tcp.keepidle=61000
isi_sysctl_cluster net.inet.tcp.keepintvl=5000
```

4. Run the following command to verify that the change was successfully added to the `/etc/mcp/override/sysctl.conf` file:

```
cat /etc/mcp/override/sysctl.conf
```

Output similar to the following is displayed:

```
<sysctl_name>=<value> #added by script
```

For example:

```
cat /etc/mcp/override/sysctl.conf
efs.bam.layout.disk_pool_global_force_spill=1 #added by script
net.inet.tcp.keepidle=61000 #added by script
net.inet.tcp.keepintvl=5000 #added by script
```

5. If you need to revert the `sysctl.conf` file to the backup version created previously:
 - a. Open an SSH connection on any node in the cluster and log on using the `root` account.
 - b. Run the following command to copy and then rename the original backup of the `sysctl.conf` file:

```
cp /etc/mcp/override/sysctl.conf.bkul /etc/mcp/override/
sysctl.conf
```

Refer to the KB Library topic: [000089232](#) for further information about configuring these parameters.

Frame loss reduction

OneFS is a scale-out, single namespace, clustered file system. To maintain coherency, OneFS implements a distributed lock manager that marshals locks across all nodes in the cluster.

When a node is added or removed from the cluster, all operations must be temporarily suspended until all existing locks are rebalanced across the resulting node set. The system must then recalculate the cluster write plan. The time required for this group change to occur depends on the size of the cluster, individual node performance, and cluster workload.

We optimized the parameters on the cluster to remove the frame loss duration.

Procedure

1. Set the parameters in the `sysctl` configuration file using the following commands:

```
declare -i COUNT MDS
BASE=10000
COUNT=$((1.01 * $BASE))
MDS=$(( $BASE * 0.75))
isi_sysctl_cluster kern.maxvnodes=$BASE
isi_sysctl_cluster kern.minvnodes=$BASE
isi_sysctl_cluster efs.lin.lock.initiator.lazy_queue_goal=
$COUNT
isi_sysctl_cluster efs.ref.initiator.lazy_queue_goal=$COUNT
isi_sysctl_cluster
efs.mds.block_lock.initiator.lazy_queue_goal=$MDS
isi_sysctl_cluster efs.bam.data.lock.initiator.lazy_queue_goal=
$MDS
```

2. Verify that the changes are logged in `sysctl.conf` file:

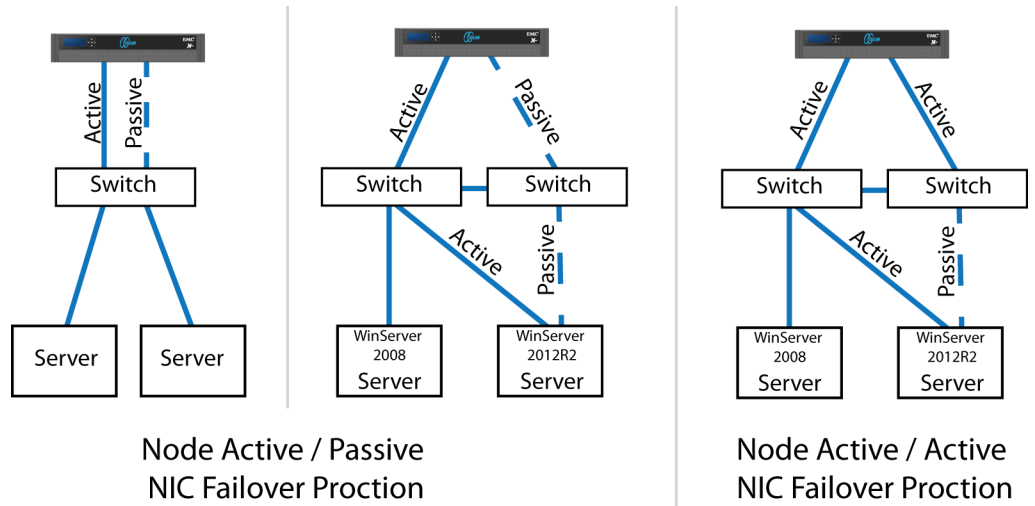
```
cat /etc/mcp/override/sysctl.conf
net.inet.tcp.keepidle=61000 #added by script
net.inet.tcp.keeptv1=5000 #added by script
kern.maxvnodes=10000 #added by script
kern.minvnodes=10000 #added by script
efs.lin.lock.initiator.lazy_queue_goal=10100 #added by script
efs.ref.initiator.lazy_queue_goal=10100 #added by script
efs.mds.block_lock.initiator.lazy_queue_goal=7500 #added by
script
efs.bam.data.lock.initiator.lazy_queue_goal=7500 #added by
script
```

Link aggregation

The active/passive configuration involves aggregating the NIC ports on the Isilon nodes for high availability. If one of the ports on the node or switch port fails, the Qognify SVR can continue writing to the Isilon share using the other port connection without affecting the recording. The SMB share continues to be accessible to the server using the passive connection port.

NIC aggregation can be used to reduce the possibility of video loss from a cable pull, NIC failure, or switch port issue. Dell EMC recommends NIC aggregation, also known as link aggregation, in an active/passive failover configuration. This method transmits all data through the master port, which is the first port in the aggregated link. If the master port is unavailable, the next active port in an aggregated link takes over.

Figure 2 Isilon Active/Passive and Active/Active configuration



I/O optimization configuration

As of OneFS 7.0.x, no changes are necessary to the I/O profiles for the directories that are used for Qognify.

Note

This setting does not require a SmartPool license.

Configuring authentication and access control

We conducted authentication and access control tests to determine the best method for shared access.

The following three tests were conducted:

Full Active Directory (recommended)

Where the Qognify server and the Isilon cluster are part of the same Windows domain.

Partial Active Directory

Where the Qognify servers are part of the Windows domain, but the Isilon cluster is administered locally.

Fully locally administered control

Where the Qognify servers and the Isilon cluster are administered locally.

Alternatives to the previous methods might exist, but the Dell EMC Surveillance Lab team does not plan to derive or support other methods.

Procedure

1. Select **Cluster Management > Access Management**.
2. Select **Access zone** and ensure that the **System access zone** has the provider status **Active Directory, Local, and File** marked with a green dot.
3. Under **Active Directory**, select **Join a domain** and add the Windows domain and appropriate users using one of the following options:

- When the Isilon cluster and Qognify are not part of the same domain, set the shares to **Run as Root**. This setting is not ideal from a security perspective.
- When the Isilon cluster and Qognify server are part of the same domain, configure the `DVM Camera` service to use the Domain account with read/write permissions to the Isilon cluster share. During the initial installation of the camera server, use the Qognify administrator account specification wizard to configure the camera service. Specify the recording location for the camera server using the full UNC path of the Isilon share.

Releases tested

The following tables list the firmware builds and software releases used for our tests.

Table 1 Firmware builds

Model	Firmware
VNX5400	VNX OE 5.33.009.5.155

Table 2 OneFS releases

Model	Firmware
X410	8.0.1.1
NL410	8.0.1.1

Table 3 Qognify releases

Release	Subrelease
Qognify	VisionHub 4.2 and NiceVision Net 3.1

Continuous Availability

Continuous Availability (CA) is a feature in OneFS 8.0 that contributes to a transparent failover during a node or NIC failure. Dell EMC recommends using CA enabled shares to minimize video loss during node or NIC failure operations.

CA describes when a node becomes inaccessible for any reason, such as administrative, failure, or infrastructure, then another node can be chosen to take its place and work can resume. CA is also known as "SMB Transparent Failover."

To improve the performance of the cluster, CA was not enabled on the shares in our testing. With CA disabled, the reconnect time for the SMB clients during the node failure and reboot scenarios is about 18 seconds by default. Optionally, this reconnect time can be reduced by using the following timeout values.

Initial RTO

Retransmit timeout (RTO) determines how many milliseconds of unacknowledged data it takes before the connection is aborted. The default timeout for Initial RTO

is 3 seconds. Use the following powershell command to set the default timeout to 2 seconds:

```
netsh int tcp set global initialRto=2000
```

TcpMaxDataRetransmissions

Determines how many times TCP retransmits an unacknowledged data segment on an existing connection. TCP retransmits data segments until they are acknowledged or until this value expires.

TCP/IP adjusts the frequency of retransmissions over time. TCP establishes an initial retransmission interval by measuring the round trip time on the connection. The interval doubles with each successive retransmission on a connection, and it is reset to the initial value when responses resume.

The default value for this parameter is 5. Reduce this value to 3 by adding the following Dword to HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Tcpip\Parameters :

```
Value Name: TcpMaxDataRetransmissions
Data Type: REG_DWORD - Number
Valid Range: 0 - 0xFFFFFFFF
Value: 3
```

Note

Make sure to backup the system before editing registry settings.

Job Engines and performance impact

During testing in the Dell EMC Surveillance Lab, we found certain job Engines that can impact the performance of the recorders.

FlexProtect and FlexProtectLin

The FlexProtect and FlexProtectLin job engines scan the file system after a device failure to ensure that all the files remain protected. FlexProtect is most efficient when used in clusters that contain only HDD drives. FlexProtectLin is most efficient when the file system metadata is stored on SSD drives.

InsightIQ®

InsightIQ takes large snapshots to provide better reporting on files that might need to be moved, which can cause performance issues in the cluster. For more information about disabling snapshots, refer to the Knowledge Base article [How to enable or disable FSAnalyze from creating a snapshot](#).

FSanalyze

FSanalyze is a job Engine that collects File System Analytics for InsightIQ reporting. If you do not need this feature, use the following command to disable FSanalyze:

```
isi job types modify fsanalyze --enabled=no
```


Impact policies

There are three impact policies: low, medium, high. Avoid changing any of the impact policies if possible, but there are cases where changes are required. Use the following information for cases where a change must be made.

To list all Impact Policy parameters

```
isi_gconfig -t job-config impact profiles
```

To list only Medium Impact Policy parameters

```
isi_gconfig -t job-config impact.profiles.medium
```

To reduce the number of workers to 50 percent

Tuning should revolve around workers per core. If the workers per core is set to 1, then the maximum number of workers in a cluster equals the maximum number of cores. For example, if the medium impact policy default is 1 worker per core, use the following command to reduce the maximum workers allowed by 50 percent:

```
isi_gconfig -t job-config impact profiles medium
workers_per_core=0.5
```

Sample Output:

```
CLuster# isi_gconfig -t job-config impact profiles medium
impact.profiles.medium.id (enum job_impact_id) = Medium
impact.profiles.medium.ionice (int) = 1
impact.profiles.medium.workers_per_core (float) = 0.5
impact.profiles.medium.min_workers_per_cluster (float) =
0.25
impact.profiles.medium.max_workers_per_storage_unit (float)
= 2
impact.profiles.medium.fixed_worker_count (float) = 0
impact.profiles.medium.max_node_load_factor (float) = 3
impact.profiles.medium.min_node_load_factor (float) = 2
impact.profiles.medium.disk_types.sata.type (enum
disk_type) = sata
impact.profiles.medium.disk_types.sata.enabled (bool) = true
impact.profiles.medium.disk_types.sata.max_client_load_kbps
(int) = 1024
```

SSD strategies

Metadata read/write acceleration writes file data to HDDs and all metadata mirrors to SSDs. This strategy accelerates metadata writes, in addition to reads, but requires about four to five times more SSD storage than metadata read acceleration. For the Isilon X410 and NL410, the Dell EMC Surveillance Lab recommends using two 1.6 TB SSDs if using metadata read/write mode.

DNS specific configuration

In our testing, we discovered that during a node or NIC failure, all the recorders in the failed node may reconnect to a single available node. In this case, round-robin does not

distribute the client connections across the available nodes and all the recorders in the failed node tried to reconnect at the exact same time.

The Microsoft DNS server caches the Node IP addresses for queries made with a time to live (TTL) of 1 second. If there are multiple recursive queries for the same DNS zone name within the same second, the DNS server responds with the same node IP for the client connection requests.

This issue can be resolved by using an alternate DNS implementation such as BIND or DNSMASQ. Another option is to use the SmartConnect service IP as the preferred DNS server and the Domain DNS server as the alternate DNS server IP.

Procedure

1. Browse to **ControlPanel > Network and sharing center > Local Area Connection**.
2. Click **Properties** in the **Local area Connection Status**.
3. Type the preferred and alternate DNS server IP addresses.

The preferred DNS is the SmartConnect service IP and the alternate DNS is the domain DNS server IP address. DNS name resolution that SmartConnect cannot answer may be sluggish in some environments.

Manually re-balancing recorders across nodes

After and activity that causes recorders for move between Isilon nodes, it is possible for the recorder to node ratio be become unbalanced. Using this procedure, a recorder may be moved from the existing node to another node in the cluster. To get the recorder to the desired node, it may take multiple iterations of the procedure.

Procedure

1. Delete the SMB sessions that allow it to reconnect to other nodes.

Type the following commands:

```
isi smb sessions list
isi smb sessions delete -f <computer name>
```

Network adapter configuration

When using the VMXNET3 driver on ESXi 4.x, 5.x or 6.x, there is significant packet loss during periods of very high traffic bursts.

To overcome this issue, the following network adapter configurations are recommended for virtual machine SVR servers.

Procedure

1. Click **Start > Control Panel > Device Manager**.
2. Right-click **vmxnet3** and click **Properties**.
3. Click the **Advanced** tab.
4. Click **Small Rx Buffers** and increase the value.

The default value is 512 and the maximum is 8192.

5. Click **Rx Ring #1 Size** and increase the value.
The default value is 1024 and the maximum is 4096.
6. Click **Tx Ring Size** and increase the value.
The default value is 1024 and the maximum is 4096.

Add a NAS drive

To add a NAS drive, create a NAS key to the registry with a string named `Supported`.

Procedure

1. **Browse to** `Computer\HKLM\Software\Wow6432node\Securitysystem\NVR`.
2. Create a key named `NAS`.
3. Add a string named `Supported` with a value data of `Y`.

Note

A hot fix is being developed to automate this procedure. If the release includes the hotfix, this procedure will not be necessary.

Set the Network storage option

To allow the SVR to reconnect to the SMB shares after a NIC or node failure, select the Network storage option in the Recorder Configuration Tool.


Procedure

1. Open the Recorder Configuration Tool.
2. Check **Network storage**.

Results

Selecting the Network storage option makes sure the recorder reconnects during node failures.

Figure 3 Recorder Configuration Tool



Snapshots and data progression

A snapshot is a point-in-time-copy (PITC) of a volume that provides fast recovery of data. Data progression moves data within a virtualized storage environment, between tiers, drive types, and between multiple RAID levels within the same tier.

Data progression is not included as part of the validation test. Dell EMC does not recommend using data progression with video surveillance workloads.

Microsoft Multipath I/O

Microsoft Multipath I/O (MPIO) is a framework that allows administrators to configure load balancing and failover processes for Fibre Channel, iSCSI, and SAS connected storage devices. Dell EMC SC Series arrays provide redundancy and failover with multiple controllers and RAID modes.

However, servers still need a way to spread the I/O load and handle internal failover from one path to the next, which is where MPIO plays an important role. Without MPIO, servers see multiple instances of the same disk device in Windows disk management.

The MPIO framework uses Device Specific Modules (DSM) to allow path configuration. Microsoft provides a built-in generic Microsoft DSM (MSDSM) for Windows Server 2008 R2 and above. This MSDSM provides the MPIO functionality for Dell EMC storage customers.

For configuration details, refer to the white paper [Dell EMC SC Series Storage: Microsoft Multipath I/O Best Practices](#).

Spare disks

Depending on the RAID level and the total number of disks in each SC Series storage array, one or more spare disks are automatically configured and used in the event of a disk failure.

Dell EMC highly recommends using spare disks as an additional level of protection in case of a disk failure. The spare disks replace the failed disk and allow the RAID set to rebuild.

CHAPTER 3

Conclusion

This chapter presents the following topics:

- [Summary](#).....30

Summary

Dell EMC performed comprehensive testing with Qognify against Dell EMC Isilon clusters. The Qognify architecture and product suite allows extreme scaling, from a few cameras to up to tens of thousands of cameras, by using Dell EMC storage.

Qognify VisionHub 4.2 and NiceVision Net 3.1 deliver a complete, end-to-end IP video surveillance VMS, that captures, records, analyzes, investigates and visualizes thousands of IP cameras. With an open platform that supports ONVIF standards, Qognify integrates with new and existing edge devices, as well as security management and access control systems.

EMC VNX arrays

The use of storage pools to create LUNs within the EMC VNX arrays greatly simplifies the configuration and increases the performance when compared to traditional block-level storage. Either iSCSI or FC can be implemented. FC performs better than iSCSI.

Dell EMC Isilon scale-out storage

Dell EMC Isilon scale-out storage is ideal for midtier and enterprise customers. An Isilon cluster is based on independent nodes working seamlessly together to present a single file system to all users.

Licensed SmartQuotas options can be configured so that each SVR view of the storage is based on the assigned quota and not the entire file system. We recommend using SmartQuotas with Qognify as a best practice.