

Surveillance

Dell EMC Storage with Avigilon Control Center

Configuration Guide

H15399

REV 1.1



AVIGILON

Copyright © 2016-2017 Dell Inc. or its subsidiaries. All rights reserved.

Published September 2016

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.

This document is not intended for audiences in China, Hong Kong, and Taiwan.

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	7
	Design concepts.....	8
	Isilon (NAS).....	8
	Large file system, small view (SmartQuotas).....	8
	Configuring SmartQuotas (recommended).....	9
	Unique share naming.....	10
	Configuring SmartConnect (optional).....	10
	SMB specific configuration.....	11
	Link aggregation.....	13
	I/O optimization configuration.....	14
	Configuring authentication and access control.....	14
	Releases tested.....	15
	Continuous Availability.....	15
	Job Engines and performance impact.....	16
	SSD strategies.....	16
	DNS specific configuration.....	17
	Manually re-balancing recorders across nodes.....	17
	Network adapter configuration.....	17
Chapter 3	Conclusion	19
	Summary.....	20
	Dell EMC Isilon scale-out storage.....	20

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 Avigilon Control Center. This document is not a replacement for the Avigilon implementation guide nor is it a replacement for the *Dell EMC Storage with Avigilon Control Center: Sizing Guide*.

Scope

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

The Dell EMC Isilon™ storage system has been tested.

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

Avigilon offers a server failover feature, which requires a hot-standby server with available storage, that was not included as part of the Avigilon Control Center validation testing. For more information about this feature, contact Avigilon.

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
- Avigilon-certified to install Avigilon Control Center services
- 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 *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](#).....8
- [Isilon \(NAS\)](#)..... 8
- [Releases tested](#)..... 15
- [Continuous Availability](#).....15
- [Job Engines and performance impact](#)..... 16
- [SSD strategies](#).....16
- [DNS specific configuration](#)..... 17
- [Manually re-balancing recorders across nodes](#).....17
- [Network adapter configuration](#)..... 17

Design concepts

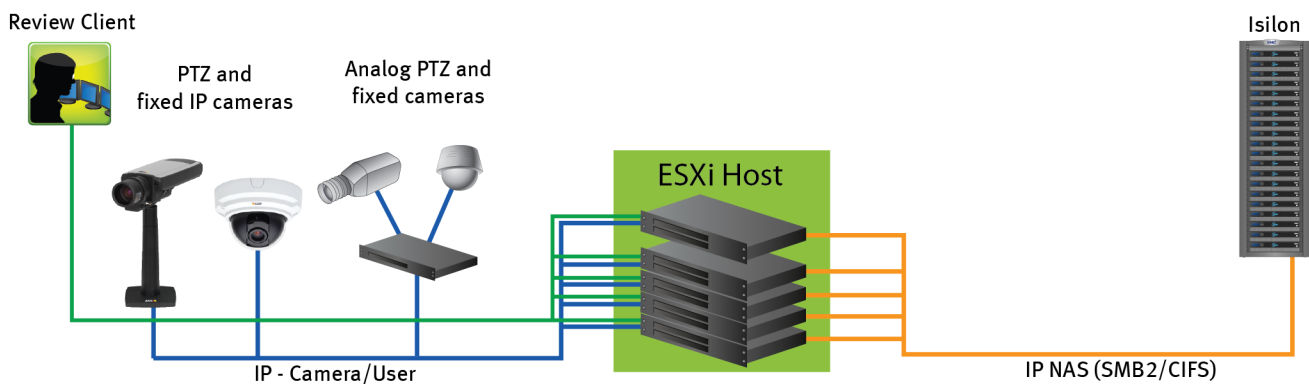
There are many design options for an Avigilon Control Center implementation. Avigilon offers many training courses related to design and implementation. These design details are beyond the scope of this paper.

The *The Avigilon Control Center Server User Guide* provides the information that you need to plan an Avigilon Video Management Software (VMS) system and components.

This guide is intended for systems integrators and architects, network IT planners, and system administrators. This guide assumes that readers know what ACC VMS does and how it works, and know how to deploy and configure Windows IP networks. These documents are available from a Avigilon partner or through the Avigilon Partner network.

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

Figure 1 Avigilon Control Center architecture



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

Large file system, small view (SmartQuotas)

Although it is possible to assign the full Isilon cluster file system to a single Avigilon Recorder, the Dell EMC best practice is to use SmartQuotas™ to segment the single Isilon file system so that each Recorder 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 Avigilon Recorder 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 Avigilon Recorder. It presents a view of available storage that is based on the assigned quota to the Recorder. SmartQuotas enables each Recorder to calculate its available disk space and react appropriately.

Configure SmartQuotas when more than one Recorder 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 Recorder 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 Recorder 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 Recorder.

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 Avigilon Recorders 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 Avigilon Recorder servers to specified nodes in the Isilon cluster. Use a unique IP address pool for video recording and Recorder read/write access.

Network throughput

Based on NIC utilization. Use of throughput requires that each Recorder 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 Recorder 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 ACC Recorder 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 ACC Recorders 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.bk1
```

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.bk1 /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.keepintvl=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 ACC Recorder 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 Avigilon.

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 tests were conducted:

Full Active Directory (recommended)

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

Partial Active Directory

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

Fully locally administered control

Where the ACC 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 Avigilon 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 ACC server are part of the same domain, configure the Avigilon Control Center service to use the Domain account with read/write permissions to the Isilon cluster share. Specify the recording location for the ACC 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 OneFS releases

Model	Firmware
NL410	8.0.0.2
HD400	8.0.0.2

Table 2 Avigilon Control Center releases

Release	Subrelease
Avigilon Control Center	5.8.2.14

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."

There are a couple of key features for this feature to work:

Transparent Failover

The "Transparent Failover" capability of SMB3 allows the connection to the shares to be maintained despite the transition the node's role (planned, unplanned). This capability allows more seamless access to the video files during most failure scenarios. Because the failover activity is not instant, to cover the short period when the client is reconnecting and reopening its files on another node, sufficient or insufficient buffering determines the video frame loss.

Witness service

Witness is service running on a second node that acts as the SMB connection witness monitoring the availability of the CA file shares. If there is a failover, the witness node has the responsibility to notify the SMB3 client that it should move its connection to a new node without waiting for TCP timeouts or DNS queries.

The Windows client starts the Witness connection. When the client makes the SMB connection to a node, the client also sends a Witness call (RPC) requesting a list of the cluster's IP addresses. The client picks the first address in the list marked as available for Witness registration. The client makes Witness calls to register for

notification. There is no load-balancing algorithm because, at the moment, Witness is only providing status change notification.

When the CA-enabled share is available cluster-wide, the SMB connection moves from one node to another node seamlessly via a reconnection.

After a connection to a node fails with non-CA enabled shares, the operating system (OS) will quickly connect to a new node. Although the OS connection to the new node exists very quickly the SMB sessions are not available immediately. SMB maintains several timeout timers that must expire before the SMB session for the failed connection is made available.

To set up continuous availability, mark the CA box when creating the share. This selection causes a bit to be set in the supported configuration mask. If the client understands and supports the option, it just happens. It is important to remember you cannot add that setting after creating a share. There are tools to support re-creating the share with the CA setting enabled but they require disconnecting all the active clients during the process.

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
```

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 an activity that causes recorders to move between Isilon nodes, it is possible for the recorder to node ratio to 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 6.0, 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.

CHAPTER 3

Conclusion

This chapter presents the following topics:

- [Summary](#).....20

Summary

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

ACC 5 is an easy to use video management software that delivers complete, end-to-end IP video surveillance that captures, records, analyzes, investigates, and visualizes. As a distributed network platform with enterprise-class reliability, ACC can efficiently capture, manage, and store high definition surveillance video while intelligently managing bandwidth and storage.

In the ACC software, servers are organized in clusters that are called sites. By organizing the system into clusters, you can control user access and system wide events through the site settings. Site settings are stored on the server, or across all servers in a multi-server system

The Isilon HD400 and NL410 clusters performed well with ACC 5. Contact the Dell EMC Surveillance Lab team with any questions regarding recommended node types.

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 recorder view of the storage is based on the assigned quota and not the entire file system. We recommend using SmartQuotas with VMS (long) as a best practice.