

Surveillance

Dell EMC Storage with Infinova 2217 Security Management System

Sizing Guide

H15952

REV 01



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

Published May 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
	Solution overview.....	6
	Scope.....	6
	Key objectives.....	7
Chapter 2	Solution components	9
	Dell EMC storage.....	10
	Storage protocols.....	10
	Infinova 2217 Security Management System.....	10
Chapter 3	Configured components	11
	Dell EMC Surveillance Lab test environment.....	12
	Isilon clustered storage system.....	12
	Data protection.....	12
	Cluster size.....	13
Chapter 4	Sizing the solution	15
	Bandwidth sizing guidelines.....	16
	Dell EMC Isilon node and cluster (SMB).....	16
Chapter 5	Testing and validation	19
	Test objectives.....	20
	Test parameters.....	20
	Tests conducted.....	20
	Video playback test	20
	Disk failure test.....	20
	NIC failure test.....	21
	NIC Failure test with NIC aggregation in Active/Passive	21
	Node poweroff test	21
	Node restart test.....	22
	Storage bandwidth and configuration test.....	23
Chapter 6	Conclusion	25
	Summary.....	26
	Dell EMC Isilon scale-out storage.....	26

CONTENTS

CHAPTER 1

Introduction

This chapter provides information on the purpose and scope of this solution:

- [Solution overview](#).....6
- [Scope](#).....6
- [Key objectives](#).....7

Solution overview

The purpose of this guide is to help you understand the benefits of using a Dell EMC storage solution with Infinova 2217 Security Management System. The solution includes both hardware and software elements for video surveillance.

Use this guide to determine the requirements for a successful Infinova 2217 Security Management System installation. The storage platforms include VMware ESXi hosts that are running Infinova 2217 Security Management System. This paper also includes information on VMware virtualization.

Scope

This guide is intended for use by internal Dell EMC sales and pre-sales personnel, and qualified Dell EMC and Infinova 2217 Security Management System partners.

The guidelines presented are for storage platform positioning and system sizing. The sizing recommendations are based on performance and storage protocol conclusions derived from Dell EMC testing.

The guidelines for sizing this video storage solution describe the use of the following storage platforms:

- Dell EMC Isilon™

These guidelines include the following design considerations:

- Architectural overview of Infinova 2217 Security Management System
- Dell EMC storage considerations for Infinova 2217 Security Management System
- Result summaries for the tests carried out by Dell EMC engineers in a VMware ESXi virtualized infrastructure

Use this guide to determine the best configuration for the following:

- Number of Infinova 2217 Security Management System Recorders
- Mix of nodes and Infinova 2217 Security Management System Recorders based on the expected bandwidth in an Isilon implementation
- Storage using Server Message Block (SMB) on Isilon systems
- Load factors related to the use of Dell EMC storage arrays in the customer's solution

Note

All performance data contained in this report was obtained in a rigorously controlled environment. Network topology and system environment variables can have significant impact on performance and stability. Follow the best practices as outlined in the *Dell EMC Storage with Infinova 2217 Security Management System: Configuration Guide* regarding network and storage array configuration. Server and network hardware can also affect performance. Performance varies depending on the specific hardware and software, and might be different from what is outlined here. Performance results will be similar if your environment uses similar hardware and network topology.

Key objectives

The configurations documented in this guide are based on tests conducted in the Dell EMC Surveillance Lab and actual production implementations.

These are the key objectives of this solution:

- Measure the sizing needs for specific system requirements so that an implementation can be correctly sized and the appropriate Dell EMC products can be matched to a customer's requirements.
- Recommend an Isilon SMB configuration.
- Calculate node maximum bandwidths.
- Recommend disk drive types.

CHAPTER 2

Solution components

This chapter provides information about storage options for video and audio data:

- [Dell EMC storage](#).....10
- [Storage protocols](#)..... 10
- [Infinova 2217 Security Management System](#)..... 10

Dell EMC storage

Dell EMC storage arrays are ideal for storing video and audio data.

This guide describes the tests for the following storage arrays:

- Isilon clusters

For our testing, we used single- and multi-node performance testing on the Isilon storage array.

Storage protocols

Dell EMC uses standard file protocols to enable users and applications to access data that is consolidated on a Dell EMC storage solution.

This guide provides information about these network protocols:

- SMB (CIFS)

Infinova 2217 Security Management System

The Infinova 2217 Security Management System V17.02.09.08 architecture can consist of a single Infinova 2217 Security Management System server or multiple Infinova 2217 Security Management System servers.

The following table describes the primary Infinova 2217 Security Management System services.

Table 1 Infinova 2217 Security Management System primary services

Service	Description
Infinova 2217 Security Management System server software	Infinova 2217 Security Management System servers are complete, high-performance network recording, and video management solutions that are designed to fit various environments, including yours. The state-of-the-art video recording, video value-added services, and analytics can turn any channel into a smart one. The Infinova 2217 Security Management System server family is fully scalable and can simultaneously manage encoders with third-party IP devices while offering a unique migration path from analog to IP.
Infinova 2217 Security Management System client software	Infinova 2217 Security Management System client software is an enhanced application suite that provides realtime event management, a network-based matrix alternative to a traditional analog matrix, investigation tools, mobile viewing options, and user management tools.

CHAPTER 3

Configured components

This chapter provides information about the components configured in this solution:

- [Dell EMC Surveillance Lab test environment](#)..... 12
- [Isilon clustered storage system](#)..... 12

Dell EMC Surveillance Lab test environment

The Dell EMC Surveillance Lab is constantly being upgraded to the most recent software releases.

In order to test this solution, the Dell EMC Surveillance Lab was configured as follows:

- 8 vCPUs
- 8 GB memory
- Network adapter type: VMXNET3 - 10 GbE
- FC Block storage for primary storage
- SMB Isilon share for secondary storage

For all the tests, the virtual CPU (vCPU), memory, and network were configured according to Infinova 2217 Security Management System best practices. The VMware vSphere configuration was in accordance with the VMware Compatibility Guide (www.vmware.com/resources/compatibility/search.php).

The Dell EMC Surveillance Lab's host hardware met and exceeded the minimum system requirements for an ESXi/ESX installation. The Infinova 2217 Security Management System Recorder VM was running on an ESXi 6.0 host using Cisco UCS B230 Blade Servers with a 20-core ESXi host at 2.2 GHz and 256 GB memory. For more information about VM configuration, see the General recommendations for storage and sizing section of the *Using EMC VNX storage with VMWare VSphere* guide.

Isilon clustered storage system

Isilon NAS was designed and developed specifically for storing, managing, and accessing digital content and other unstructured data.

An Isilon clustered storage system is composed of three or more nodes. Each node is a self-contained, rack-mountable device that contains industry-standard hardware such as disk drives, CPUs, memory, and network interfaces. These nodes are integrated with the proprietary Isilon OneFS™ operating system, which is a distributed networked file system that unifies a cluster of nodes into a single shared resource.

Data protection

OneFS does not rely on hardware-based RAID for data protection. The Isilon system uses the Reed-Solomon algorithm for N+M protection with Forward Error Correction (FEC).

Protection is applied at the file level, enabling the cluster to recover data quickly and efficiently. Nodes, directories, and other metadata are protected at the same or a higher level as the data blocks they reference. Since all data, metadata, and FEC blocks are spread across multiple nodes, dedicated parity drives are not required. For more information about Isilon data protection, see *Dell EMC Isilon OneFS: A Technical Overview*.

Although cluster sizes as small as three nodes are possible, for surveillance applications we recommend a minimum of five nodes. Sizing calculations need to include a minimum free space calculation for proper cluster sizing. We recommend a cluster size that enables a node to be removed while retaining a minimum of 10 percent free space in the remaining capacity. This cluster size ensures that node removal and node failures have minimal or no impact on video ingestion.

The Isilon sizing tool provides an accurate calculation. You can find this tool at <https://isilon-sizing-tool.herokuapp.com>. Other sizing tools from video management software (VMS) and camera vendors may also be used for sizing the necessary bandwidth and storage capacity.

Isilon protection with OneFS

New or upgraded clusters, starting with OneFS 7.2, provide a data protection level that meets Dell EMC Isilon guidelines for mean time to data loss (MTTDL) for large capacity nodes. Current releases of OneFS offer a new protection option, 3d:1n1d, which means the cluster can survive three simultaneous disk failures or one entire node failure plus one disk. OneFS also provides an option that continually evaluates the cluster and sends an alert if the cluster falls below the suggested protection level.

Cluster size

We recommend a minimum cluster size of five nodes, even if you are not writing to all of them. For example, if you are implementing a four-node Recorder solution, implement a five-node cluster. This also meets the recommended best practices for data protection.

To estimate the ideal number of nodes in a cluster, you need to consider cluster bandwidth and capacity.

Sizing by bandwidth

We recommend a cluster size with one or more additional nodes than calculated in bandwidth sizing. This ensures that failover of a node allows for redistribution of NAS connections and avoids any frame loss.

Sizing by aggregate capacity

We recommend a cluster size with enough usable capacity to handle 110 percent of the calculated space requirement, with a minimum added capacity of one full node plus 10 percent. The values are based on camera bit rate.

The Isilon sizing tool can use both the sizing by bandwidth and sizing by aggregate capacity methods when calculating ideal cluster size.

Configured components

CHAPTER 4

Sizing the solution

This chapter provides information to enable you to quickly determine the correct storage array based on your customer's bandwidth requirements:

- [Bandwidth sizing guidelines](#)..... 16
- [Dell EMC Isilon node and cluster \(SMB\)](#)..... 16

Bandwidth sizing guidelines

All solution tests were performed in a lab environment. The storage system, cameras, and VLANs in the lab environment were dedicated to these tests.

Connections to the storage system under test conditions were restricted to Infinova 2217 Security Management System Recorder, monitoring, and web management stations. Expect some variance between the lab results and a production environment.

Dell EMC Isilon node and cluster (SMB)

The test results are based on a model in which the constant-bandwidth surveillance video traffic remained unaffected during a single node maintenance cycle, disk rebuild, or non-disruptive upgrade.

We used 10 Gigabit Ethernet (GbE) interfaces. We performed all tests with a per camera bandwidth of 4 Mb/s, so a single Recorder that handles 32 MB/s can support 64 such cameras.

We performed all tests with node or drive failures in place in the cluster (for example, with Isilon FlexProtect™ running) to ensure a worst-case scenario for all sizing parameters. Maximum per recorder bandwidth may vary based on the configuration of the SVR server used.

The following table provides bandwidth-sizing guidelines based on our test results.

Table 2 Dell EMC Isilon node and cluster (SMB) test results

Cluster	OneFS version	SVRs per node	Bandwidth (MB/s)		Drive size	Maximum cluster RAW
			Per SVR	Per node		
HD400	8.0.1.1	1	30	30	6 TB	30.2 PB
	8.0.1.1	2	30	60	6 TB	
	8.0.1.1	3	30	90	6 TB	

Note

All disk drives are NL-SAS 7200 RPM unless otherwise noted.

This guide provides details on the total load that was tested in the Dell EMC Surveillance Lab. However, the independent software vendor (ISV) should provide the actual server specification. The test results in this guide set a server bandwidth specification that is based on our lab environment, which can be used in the event the ISV does not provide these specifications.

Note

Isilon SMB shares were used for secondary storage.

SSD drive 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. If using

the metadata read/write mode for the Isilon X410 and NL410, the Dell EMC Surveillance Lab recommends using two 1.6 TB SSDs per node.

CHAPTER 5

Testing and validation

This chapter describes the testing used to validate this solution.

- [Test objectives](#).....20
- [Test parameters](#).....20
- [Tests conducted](#)..... 20
- [Storage bandwidth and configuration test](#)..... 23

Test objectives

Many factors must be considered when designing your solution.

The Dell EMC Surveillance Lab tests focus on storage-related factors with the following objectives:

- Determine the bandwidth for various Dell EMC storage clusters using SMB.
- Determine the best configuration parameters for Isilon storage options.
- Determine best video storage performance requirements for use with Isilon scale-out storage clusters.
- Determine the maximum bandwidth with multiple Recorders.
- Determine all factors with a lab-controlled failure, such as rebuilding disks, removing a node, or network path failures.

Test parameters

All test parameters and scenarios reflect standard production behavior for Infinova 2217 Security Management System under storage-intensive conditions, including typical storage functions and failures. We followed best practices for recovery and break-fix issues for normal situations that might arise in a standard production environment.

We used the following parameters to perform the tests:

- The IP network (Layer 2) is a flat, high-availability network with plenty of capacity, which enabled us to focus on the products we were testing.
- All tests assumed uniform distribution of bandwidth from the Infinova 2217 Security Management System Recorder.

Tests conducted

We ran tests with the SmartConnect™ configuration in place and the SMB shares were mounted using the SmartConnect zone name.

Video playback test

As video is being written to the storage, we recalled or reviewed the video at a rate equal to 20 percent of the write rate. We ran tests with the SmartConnect™ configuration in place and the SMB shares were mounted using the SmartConnect zone name.

The review did not affect the write rate, video quality, or result in dropped video.

Disk failure test

A single disk failure is the most common failure affecting storage systems today. When a disk fails, that disk is removed and replaced. The replacement disk is then reconstructed.

The Isilon cluster was protected using a +2 protection scheme that allows for two simultaneous disk failures. In our test, we failed and recovered two disks. The

SmartFail process started and the CPU utilization of the node increased with no observed effect to the write streams.

NIC failure test

We performed the hard NIC failure test by removing one NIC cable from the active node that was involved in active recording. After the NIC failure, writing to the same node failed. When the network fails, the server must recognize the failure, then it must establish a new connection. Also, when the network fails TCP socket connections are left open and remain open on the cluster until Isilon's OneFS forces them closed, which allows the server to continue writing.

We can force the open TCP sockets to close for a duration of less than 2 minutes by reducing the `TCP keep idle` and `TCP keep interval` timeout to the optimum values recommended by Isilon Engineering.

To reduce the video loss duration due to the `TCP Socket Open` condition, set the persistent values in the `sysctl.config` file as follows to reduce the impact duration time significantly:

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

Refer to the KB article: 000089232 for further information about how to configure these parameters.

Note

NIC failure impact can be overcome by using NIC aggregation in Active/Passive Failure aggregation mode, which is explained in the next test case. Connectivity to the nodes that are not affected by the network outage continues to be available throughout the test scenario and no impact was observed.

NIC Failure test with NIC aggregation in Active/Passive

We ran a hard NIC failure test with Active/Passive aggregation by removing the active NIC port cable. After the network failure, writing to the same node continued and the NIC that was passive was immediately changed to the active NIC. The NIC failure caused no apparent loss.

Note

NIC aggregation in Active/Passive mode remedies only a network disconnection/NIC failure that happens on the Isilon node or the corresponding switch port where it is connected.

Node poweroff test

We simulate an unexpected single node hard failure, which causes the servers that were writing to that node to reconnect to a new node.

In our tests, the servers on the failed node reconnected to a new node, but did not start writing again for an aggregate (reconnect and start writing) duration of up to 52 seconds while waiting for writing to the SMB share to be re-started.

The second issue is that the removal or addition of a node causes an interrupt to the cluster. Therefore, video servers writing to the other nodes might experience a short interruption. The duration of the interruption can be reduced by modifying the OneFS environment variables.

The following changes are required to modify the remove or add node interruption:

```
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
```

⚠ WARNING

- If running a mixed workload, these changes can adversely affect the other workloads that might be present on the cluster.
- During the node failure test, writing to both primary and secondary storage stopped working for about 30 to 40 seconds, until the recorders reconnected to smart connect.
- We also tested shutting down secondary storage (Isilon cluster) completely. In this test writing to the primary storage stopped and did not recover. This issue can be addressed by rebooting the VMS recorders. After the reboot, writing to the primary storage continues in the absence of a secondary storage.

Node restart test

In our tests, we observed that a manual restart of a node would induce a video loss of 30-40 seconds.

This loss can be addressed by moving the SMB client connections on that node across to other nodes before the restart:

1. Suspend the node SmartConnect: Run the following command on the node that must be rebooted.
For example, for node 3 with groupnet0, Subnet0 and pool0, run the following command:

```
isi network pools sc-suspend-nodes groupnet0.subnet0.pool0 3
```

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

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

Note

After the node restart is completed, you can enable SmartConnect using the command `isi network pools sc-resume-nodes groupnet0.subnet0.pool0 3`.

Storage bandwidth and configuration test

The storage bandwidth test evaluated video storage and applications with a number of different Dell EMC storage systems. Additional tests evaluated ESXi host hardware in relationship to vCPU settings and the resulting bandwidths.

These tests assumed that Infinova 2217 Security Management System Recorder was configured as described by Infinova 2217 Security Management System best practices and operated within the recommended bandwidth, camera count, and other Infinova 2217 Security Management System maximum requirements.

Procedure

1. Configured video storage for a Dell EMC storage system.
2. Set up camera simulators (traffic generators) to produce a traffic load to each Infinova 2217 Security Management System Recorder at the desired bandwidth.
3. Verified that Recording is set to use H.264 and the max resolution is set to 25 FPS (1920x1080).
4. Evaluated the network and video storage to ensure an error-free environment at the induced bandwidth.
5. Introduced storage device errors including:
 - Disk failures and rebuilds on Isilon nodes
 - Initiation of Isilon node failures and recoveries
 - Initiation of Isilon node removals (downsizing a cluster)
 - Initiation of Isilon node SmartFail
6. Captured the storage system and host statistics.
7. Based on the test results:
 - If no issues were detected, incremented the bandwidth.
 - If issues were detected, decreased the bandwidth.

This procedure was repeated until the maximum error-free bandwidth was determined.

Testing and validation

CHAPTER 6

Conclusion

This chapter summarizes the testing for this solution:

- [Summary](#).....26

Summary

Dell EMC performed comprehensive testing with Infinova 2217 Security Management System against Dell EMC Isilon clusters. The Infinova 2217 Security Management System architecture and product suite allows extreme scaling, from a few cameras to up to tens of thousands of cameras, by using Dell EMC storage.

Infinova 2217 Security Management System V17.02.09.08 delivers complete, end-to-end IP video surveillance that captures, records, analyzes, investigates and visualizes. With an open platform that supports ONVIF standards, Infinova 2217 Security Management System integrates with new and existing edge devices, as well as security management and access control systems.

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 Infinova 2217 Security Management System as a best practice.