



Physical Security EMC Storage with ISS SecurOS

Version 1.0

Sizing Guide

H14192



Copyright © 2015 EMC Corporation. All rights reserved. Published in USA.

Published June, 2015

EMC 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. EMC Corporation 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 EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

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

CONTENTS

Chapter 1	Introduction	5
	Solution overview.....	6
	Scope.....	6
	Key recommendations.....	6
Chapter 2	Configured components	9
	Isilon clustered storage system.....	10
	Data protection.....	10
	Isilon protection with OneFS 7.2.....	11
	Cluster size.....	11
Chapter 3	Solution components	13
	EMC storage.....	14
	Storage protocols.....	14
	ISS SecurOS Enterprise.....	14
Chapter 4	Sizing the solution	17
	Bandwidth sizing guidelines.....	18
	EMC VNX	18
	EMC Isilon node and cluster (SMB2).....	19
Chapter 5	Testing and validation	21
	Test objectives.....	22
	Test parameters.....	22
	Tests conducted.....	22
Chapter 6	Conclusion	25
	Summary.....	26
	EMC Isilon scale-out storage.....	26
	EMC VNX arrays.....	26
	EMC VNX-VSS100 arrays.....	26
Index		27

CONTENTS

CHAPTER 1

Introduction

This document provides sizing guidelines for storage arrays and storage clusters, with sizing based on performance requirements and storage protocols.

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

Solution overview

Video surveillance is a highly competitive market, not only for Video Management Software (VMS) providers, but also for hardware and value-added companies such as EMC.

The purpose of this guide is to help you better understand the benefits of using an EMC storage solution for video surveillance that includes both hardware and software elements.

The storage platforms include physical hosts that are running SecurOS Enterprise Video Surveillance Manager.

Scope

This guide provides guidelines for storage platform positioning and sizing for system design. The sizing recommendations are based on performance and storage protocol conclusions derived from EMC testing.

This guide is intended for use by internal EMC sales and pre-sales personnel, and qualified EMC and ISS partners.

This guide provides guidelines for sizing ISS SecurOS using EMC storage systems for video storage including:

- EMC Isilon®
- EMC VNX®

It includes the following design considerations:

- Architectural overview of ISS SecurOS
- EMC storage considerations for ISS SecurOS
- Result summaries for the tests carried out by EMC engineers

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

- Number of Video Servers
- Mix of nodes and Video Servers based on the expected bandwidth in an Isilon implementation
- Storage using Fibre Channel (FC) and Internet SCSI (iSCSI) on the VNX
- Storage using Server Message Block (SMB 2.0) on the Isilon
- Load factors related to designing EMC storage arrays into the customer's solution

All performance data contained in this report was obtained in a rigorously controlled environment. Performance varies depending on your specific hardware and software and might be different from what is outlined here.

Key recommendations

The configurations documented in this paper are based on tests that were conducted in the EMC Physical Security lab, production implementations, or a combination of the two.

Following are the key recommendations in this paper, which are based on the results of these tests:

- Measure sizing needs for EMC storage options so that an implementation can be correctly sized and the right EMC product(s) can be matched to the customer's requirements
- Determine EMC VNX bandwidth for each logical unit number (LUN)
- Determine disk pool and LUN configurations
- Recommend Isilon SMB configuration
- Determine Isilon bandwidth for each ISS Video Server
- Calculate array or node maximum bandwidths
- Recommend disk drive types

CHAPTER 2

Configured components

This section provides information about the components configured in this solution.

- [Isilon clustered storage system](#)..... 10
- [Data protection](#)..... 10
- [Cluster size](#)..... 11

Isilon clustered storage system

Isilon scale-out network-attached storage (NAS) was designed and developed specifically to address the needs of 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 chips, 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.

Network interface card (NIC) aggregation can be used to reduce the possibility of video loss from a cable pull, NIC failure, or switch port issue. 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.

Data protection

Isilon OneFS does not rely on hardware-based RAID for data protection.

The Isilon system uses the Reed Solomon algorithm for N+M protection. In the Isilon N+M data protection model, N represents the number of nodes, and M represents the number of simultaneous node, drive, or node and drive failures that the cluster can withstand without incurring data loss. N must be larger than M. OneFS supports N+1, N+2, N+3, and N+4 data protection schemes, and up to 8x mirroring. OneFS also supports several hybrid protection schemes. These include N+2:1 and N+3:1, which protect against two drive failures or one node failure, and three drive failures or one node failure, respectively.

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 higher level as the data blocks they reference. Since all data, metadata, and forward error correction (FEC) blocks are spread across multiple nodes, there is no requirement for dedicated parity drives.

The following best practices are based on a five-node minimum cluster size. It is possible to use cluster sizes as small as a three-node cluster, but it is not recommended.

- Isilon recommends a +2:1 protection level for five-node clusters. Larger node-count clusters have more disks, which increases the possibility of multiple disk failures. For larger clusters, consult the EMC Isilon team for appropriate protection schemes: N+2:1, N+2, N+3, or N+4.
- 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 a more accurate calculation. You can find this tool at <https://isilon-sizing-tool.herokuapp.com>. Other sizing tools from VMS and camera vendors may also be used for sizing the bandwidth and storage capacity that is needed.

During our testing, we discovered that there might be some video loss when adding or removing a node from the cluster. Isilon 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 that is required for this group change to occur depends on the size of the cluster, individual node performance, and cluster workload.

Isilon protection with OneFS 7.2

On new installations, OneFS 7.2 provides a data protection level that meets EMC Isilon guidelines for mean time to data loss (MTTDL) for large capacity nodes. OneFS offers a new protection option, +3d:1n1d. OneFS 7.2 also provides an option where it 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-Video Server 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 can base the calculation on either bandwidth or capacity.

Sizing by bandwidth

We recommend a cluster size with one or more additional nodes than is calculated in bandwidth sizing to ensure that failover of a node allows for redistribution of NAS connections and to avoid 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 (values based on camera bit rate).

Configured components

CHAPTER 3

Solution components

This section provides information about storage options for video and audio data.

- [EMC storage](#)..... 14
- [Storage protocols](#)..... 14
- [ISS SecurOS Enterprise](#)..... 14

EMC storage

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

This sizing guide describes the tests for each of these arrays:

- Isilon clusters
- VNX arrays

For our testing, we used both single and dual storage processors for the full range of VNX storage arrays, and single, multi node performance testing on the Isilon storage array.

Storage protocols

EMC supports many storage protocols.

This guide provides information about the following network protocols:

- FC
- iSCSI
- SMB (CIFS)

ISS SecurOS Enterprise

An ISS SecurOS Enterprise installation can consist of a single video server or multiple servers in a hierarchical structure. You can configure SecurOS to manage a few cameras or thousands of cameras.

The following table describes three primary SecurOS services.

Table 1 SecurOS primary services

Service	Description
ISS SecurOS Video Server	The Video Server is a computer that is used to connect cameras (IP and/or CCTV analog cameras), from which to receive video. In addition to cameras, this installation type enables the installation or connection of various security devices (fire alarm system controllers, access control subsystem controllers, among others) and ISS analytics modules (for example, SecurOS Auto Module). The Video Server installation also includes all the functionalities of the Administrator Workstation and Operator Workstation installation types.
ISS SecurOS Operator Workstation	The Operator Workstation is a client computer in the security network designed for remote viewing of video from surveillance cameras, listening to audio, and managing various devices such as PTZ devices, doors, fire alarms and other subsystems.
ISS SecurOS Administrator Workstation	The Administrator Workstation is a client computer similar to the Operator Workstation but with an additional administrator interface for configuring the security network. From the Administrator Workstation you can perform the following administrator operations: <ul style="list-style-type: none"> • Create and configure system objects • Delete existing objects

Table 1 SecurOS primary services (continued)

Service	Description
	<ul style="list-style-type: none"><li data-bbox="679 331 1337 363">• Change the parameters of the created and operating objects

CHAPTER 4

Sizing the solution

You can use the information in this section to quickly determine the correct storage array based on your customer's bandwidth requirements.

- [Bandwidth sizing guidelines](#)..... 18
- [EMC VNX](#) 18
- [EMC Isilon node and cluster \(SMB2\)](#)..... 19

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 SecurOS Video Server, monitoring, and web management stations. Expect some variance between the lab results and a production environment.

EMC VNX

We conducted the functional tests to determine how ISS works with VNX storage arrays.

The test results shown in the following table are based on a conservative model to ensure that the constant-bandwidth video traffic is unaffected during a single storage pool (SP) maintenance cycle, disk rebuild, or similar performance-intensive events.

The VNX 5300 and VNX-VSS100 can exceed 500 MB/s in an environment without induced errors.

Table 2 EMC VNX/VNXe storage array results

Storage Protocol	Model	Single SP Bandwidth (MB/s)		Storage Pools			RAID	Disks		
		Array/Pool	LUN	Pools	LUNs/pool	Disks/pool		Size	RPM	Type
iSCSI	VNX5300	60	60	1	1	72	5	4 TB	7200	NL-SAS
iSCSI	VNX-VSS100	60	60	1	1	72	5	4 TB	7200	NL-SAS
iSCSI	VNX5200	60	60	1	1	81	5	4 TB	7200	NL-SAS
iSCSI	VNX5200	300*	30	1	5	81	5	4 TB	7200	NL-SAS

* This bandwidth is the maximum tested per SP on the VNX system and not the array maximum.

The following table lists the firmware builds used in these tests. Expect equal or better results if you use more recent firmware or newer models of VNX storage.

Table 3 EMC VNX Firmware builds

Array	Firmware
VNX5300	VNX OE 5.32.000.5.206
VNX5200	VNX OE 5.33.000.5.035

EMC Isilon node and cluster (SMB2)

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 similar performance-intensive operation.

We used 10 Gigabit Ethernet (10 GbE) interfaces with no more than four SMB connections per interface. A 1 GbE interface can accommodate up to a single Video Server connection at the maximum ISS SecurOS-supported values.

We performed all tests with a per-camera bandwidth of 7 Mb/s, so a single Video Server that handles 65 MB/s can support 75 such cameras. If the planned deployment requires many more cameras per Video Server, for example 200 low-resolution cameras, the numbers in the following table are valid with a high confidence level for bandwidth performance for up to 2,000 cameras.

We performed all tests with node or drive failures in place in the cluster (for example, with FlexProtect running) to ensure a worst-case scenario for all sizing parameters.

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

Table 4 EMC Isilon node and cluster (SMB) test results

Array	OneFS version	Video Servers per node	Bandwidth (MB/s)			Cluster size	Nodes written	Disk Drives		
			Total cluster	Per node	Per host			Size	Type	RPM
X410	7.2.0.0	1	63.3	63.3	63	4	1	4 TB	SATA	7,200
X410	7.2.0.0	2	128	128	64	4	1	4 TB	SATA	7,200
X410	7.2.0.0	4	256	256	64	4	1	4 TB	SATA	7,200
X410	7.2.0.0	3	386	193	64	4	2	4 TB	SATA	7,200

CHAPTER 5

Testing and validation

This section describes the testing used to validate this solution.

- [Test objectives](#) 22
- [Test parameters](#) 22
- [Tests conducted](#) 22

Test objectives

Many factors should be taken into account when designing your solution.

The EMC lab tests focus on storage-related factors with the following objectives:

- Determine the bandwidth for various EMC storage arrays and clusters using SMB, FC, and iSCSI.
- Determine the configuration parameters for Isilon and VNX storage options.
- Determine optimal video storage performance requirements for use with Isilon scale-out storage clusters and VNX storage arrays.
- Determine the maximum bandwidth with multiple Video Servers.

Test parameters

All test parameters and scenarios reflect standard production behavior for ISS SecurOS 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:

- Tests were conducted by loading each recorder at the maximum sustainable bandwidth of 64 MB/s, with 7Mbps camera bandwidths were used.
- The IP network (Layer 2) is a flat, high availability network with plenty of capacity, which allowed us to focus on the products we were testing.
- All tests assumed uniform distribution of bandwidth from the ISS Video Server.

Tests conducted

We ran tests with the Smartconnect configuration in place and the SMB shares were mounted using Smartconnect zone name:

Video playback at 20%

As video is being written to the storage, we recall or review the video at a rate equal to 20% of the write rate.

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 affect 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 needs to 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 two 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 paramters.

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 which 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 did 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

To simulate an unexpected single node hard failure, we held down the power button until the node powered off. This 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 code makes the group changes to the cluster that reduce the interruption from 30 seconds to 1 minute to 9 seconds or less. The changes affect the "lazy queue" and other cache related operations on each node.

These 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

These changes can adversely affect other workloads that might be present on the cluster.

CHAPTER 6

Conclusion

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

Summary

We performed comprehensive testing with ISS SecurOS against many EMC VNX arrays and EMC Isilon clusters.

The ISS SecurOS architecture and product suite allows extreme scaling, from a few cameras to up to tens of thousands of cameras, by using EMC storage.

EMC Isilon scale-out storage

EMC Isilon scale-out storage is ideal for mid-tier 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 Video Server view of the storage is based on the assigned quota and not the entire file system. In our tests, we found this feature to be necessary to guarantee a successful disk rebuild and for various node removal tests. We recommend using SmartQuotas with ISS SecurOS as a best practice.

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.

EMC VNX-VSS100 arrays

The VNX Video Surveillance Storage (VSS) is a storage solution that is purpose built to meet the unique demands of the video surveillance environment.

We found that this high availability, low-cost array performs comparably to other arrays in the VNX family.