

Dell EMC Storage with Milestone XProtect Corporate

Surveillance

September 2018

H14502.8

Sizing Guide

Abstract

This guide provides guidelines for sizing the Dell EMC storage arrays and storage clusters with Milestone XProtect Corporate 2018. The sizing recommendations are based on performance and storage protocol conclusions derived from Dell EMC Surveillance Lab testing.

Dell EMC Solutions



Copyright © 2014-2018 Dell Inc. or its subsidiaries. All rights reserved.

Published September 2018

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
	Purpose.....	6
	Scope.....	6
	Key objectives.....	7
Chapter 2	Solution components	9
	Dell EMC storage.....	10
	Storage protocols.....	10
	Milestone XProtect Corporate architecture.....	10
	Releases tested.....	13
Chapter 3	Configured components	15
	Dell EMC Surveillance Lab test environment.....	16
	VMware ESXi environment.....	16
	Unity and SC series arrays.....	17
	Isilon clustered storage system.....	17
	Dell EMC ECS Object Storage	17
	Dell EMC CIFS-ECS	18
	Dell Embedded Box PCs.....	18
	Dell EMC PowerEdge servers.....	18
Chapter 4	Sizing the solution	19
	Sizing guidelines.....	20
	Volume limits.....	20
	Live DB only sizing.....	20
	Archive DB sizing.....	21
	Dell EMC ECS	24
	Dell Embedded Box PCs.....	24
	Dell EMC PowerEdge server storage.....	25
Chapter 5	Testing and validation	27
	Test objectives.....	28
	Test parameters.....	28
	Tests conducted.....	28
	Video playback test	28
	Disk failure test.....	28
	Disk failure test with ECS.....	29
	NIC failure test.....	29
	NIC failure test with ECS.....	30
	NIC Failure test with NIC aggregation in Active/Passive	30
	Node poweroff test	31
	Node poweroff test with ECS	31
	Node reboot test with ECS.....	32
Chapter 6	Conclusion	33
	Summary.....	34

CONTENTS

Dell EMC Unity and SC series arrays..... 34
Dell EMC Isilon scale-out storage..... 34

CHAPTER 1

Introduction

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

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

Solution overview

Milestone XProtect Corporate is a tiered solution that works well with Dell EMC Unity™ or EMC VNX™ family storage arrays as well as Dell EMC Isilon™ scale-out storage. The first tier of storage, Live DB, can accommodate stored video for the retention period prior to being moved to the second tier, of storage, Archive DB or deleted. The best practice retention time is between 2 and 24 hours. The second storage tier is Archive DB, which can accommodate long video retention cycles prior to being deleted from the second tier of storage.

Live DB requires block storage, such as Dell EMC Unity or EMC VNX family arrays using iSCSI or Fibre Channel (FC), or storage local to the server. In a virtualized environment, the Unity or VNX family arrays can serve a dual purpose by providing storage for the LiveDB and VMware datastores. The Archive DB used as the secondary storage for video and surveillance data can be stored either on an Isilon scale-out storage cluster or on a unique Unity or VNX family array.

While the Milestone XProtect Corporate tiered storage solution can be deployed within a site, depending on the requirements XProtect Corporate can also provide a solution for distributed to central site architectures.

Purpose

This guide provides guidelines for sizing the Dell EMC storage arrays and storage clusters. The sizing recommendations are based on performance and storage protocol conclusions derived from Dell EMC testing.

Use this guide to determine the requirements for a successful Milestone XProtect Corporate 2018 installation, Milestone storage tiering options, and storage-specific configuration requirements.

Scope

This guide is intended for use by internal Dell EMC sales and pre-sales personnel, and qualified Dell EMC and Milestone 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™
- Dell EMC Unity™
- Dell EMC SC series™
- EMC VNX™
- EMC VSS™

These guidelines include the following design considerations:

- Architectural overview of Milestone XProtect Corporate
- Dell EMC storage considerations for Milestone XProtect Corporate
- 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 Milestone Recorders
- Mix of nodes and Milestone Recorders based on the expected bandwidth in an Isilon implementation
- Storage using Fibre Channel (FC) and Internet SCSI (iSCSI) on VNX systems
- 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

Although this document outlines some configuration parameters, it is not intended as a configuration guide. Verify all configurations with the appropriate vendor documentation or representative.

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 Milestone XProtect Corporate: 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 VNX cache configuration.
- Determine the VNX and VSS LUN bandwidth within the storage pool.
- Recommend an Isilon SMB configuration.
- Calculate array or node maximum bandwidths.
- Recommend disk drive types.
- Maximum Milestone XProtect recording server bandwidth with the Live database (Live DB) on VNX, VNXe, and Isilon storage.
- Maximum Milestone XProtect recording server bandwidth with the Live DB and Archive database (Archive DB) implemented on the same VNX5800™ storage array.
- Maximum bandwidth per Isilon node for the Archive DB.
- Confirm the previous test results with lab controlled failures, such as disabled storage processors, disk rebuilds, node removals, and network path failures.

CHAPTER 2

Solution components

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

- [Dell EMC storage](#).....10
- [Storage protocols](#)..... 10
- [Milestone XProtect Corporate architecture](#)..... 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
- Unity arrays
- SC series arrays
- VNX arrays
- VSS arrays

For our testing, we used both single and dual storage processors for the full range of Unity and SC series storage arrays and 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:

- FC
- iSCSI
- SMB (CIFS)

Milestone XProtect Corporate architecture

Milestone XProtect Corporate uses a distributed architecture with a management server as the core server. The management server can be centrally located or distributed to multiple sites and connected using the Milestone Federated Architecture. The number of recording servers is unlimited.

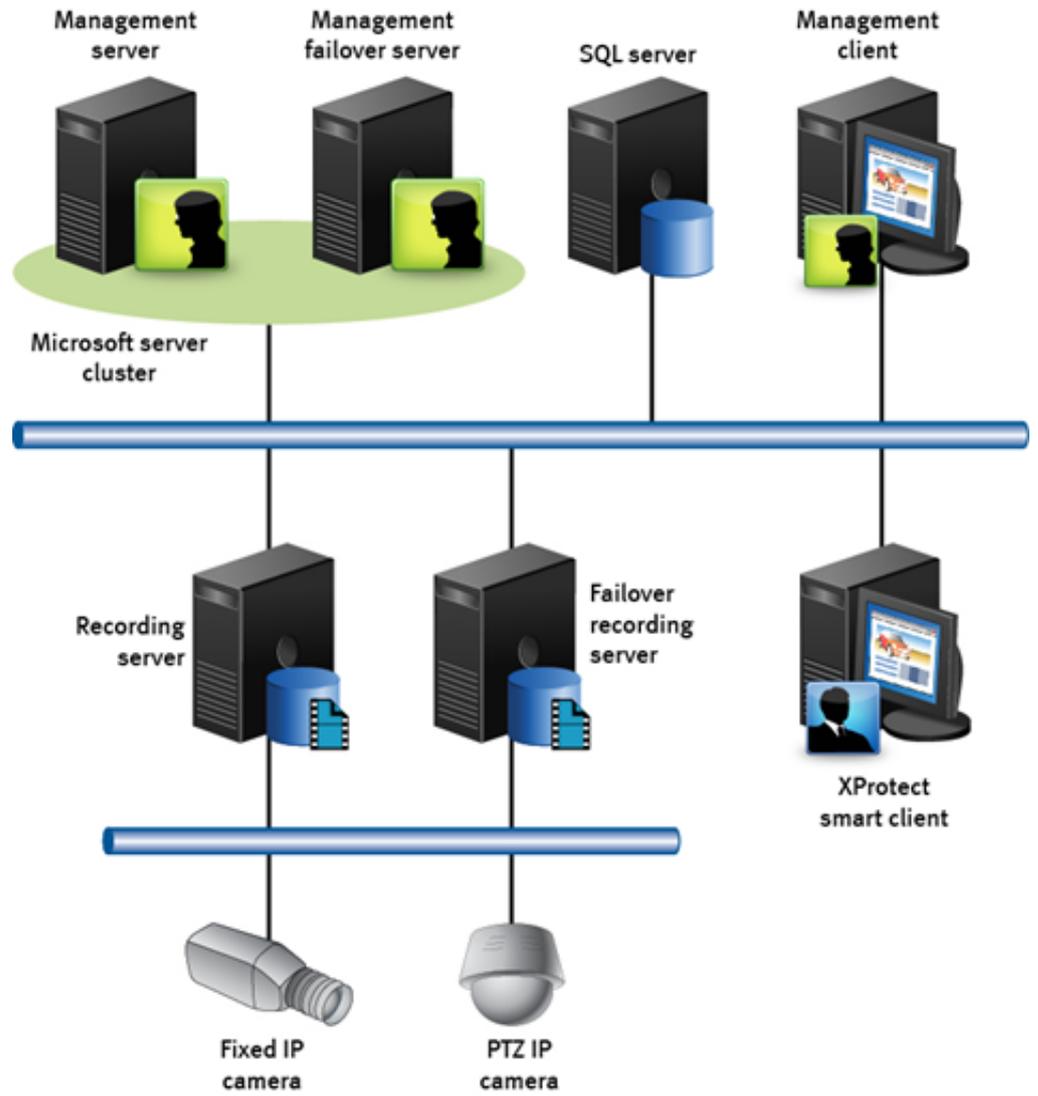
The following table lists XProtect servers, services, and their functions.

XProtect Corporate Server/ Service	Functions
Smart Client	Full-featured remote client, which provides these daily functions: <ul style="list-style-type: none"> • Simultaneous live view and playback of 100 cameras • Intelligent Pan Tilt Zoom (PTZ) camera control • Advanced search capabilities • Export of evidence material
Remote Client	Provides live view and playback of up to 16 cameras and performs most daily operations
Matrix	Allows up to four live video streams to be sent to an XProtect Corporate Smart Client PC

XProtect Corporate Server/ Service	Functions
Recording/failover server	Provides the following functions: <ul style="list-style-type: none"> • Storage and retrieval of video and audio from MJPEG, MPEG4, MxPEG, and H264 devices • Standby for a single or a group of recording servers, when configured as a failover server • Edge Storage capability, which allows cameras to write to an Edge Storage device if the recording server is unreachable • Processing events, alerts, and actions
Management server	The Management Application is XProtect's user interface to the management server and provides the following functions: <ul style="list-style-type: none"> • Managing recording servers, users, and devices • System configuration wizards, automated device discovery, smart bulk configuration, event/alarm configuration, and management of user access privileges • Multi-stage storage schemes, which enable video migrations from primary storage (Live DB) to secondary storage (Archive DB) • Hosting and controlling access from XProtect Corporate clients • Logging

The following figure shows a simple Milestone XProtect Corporate architecture. You can scale the system by expanding the number of servers in each site, and you can combine many sites into a federated architecture.

Figure 1 Milestone XProtect Corporate architecture



Releases tested

The following table lists the Milestone XProtect Corporate releases used for our tests.

Table 1 Milestone XProtect Corporate releases

Release	Subrelease
Milestone XProtect Corporate	2016, 2017, 2018

CHAPTER 3

Configured components

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

- [Dell EMC Surveillance Lab test environment](#)..... 16
- [Unity and SC series arrays](#)..... 17
- [Isilon clustered storage system](#)..... 17
- [Dell EMC ECS Object Storage](#) 17
- [Dell EMC CIFS-ECS](#) 18
- [Dell Embedded Box PCs](#)..... 18
- [Dell EMC PowerEdge servers](#)..... 18

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:

Virtualized environment:

- 8vCPUs
- 16 GB memory
- Network adapter type: VMXNET3 (GbE and 10 GbE), E1000, or VMXNET2 (GbE only)
- Isolated VLAN for storage (if not FC)

Physical/Baremetal minimum environment:

- 8 Cores
- 32 GB memory

Network environment:

- Network adapter type: 10 GbE
- Camera user VLAN
- Storage VLAN

All storage and server tests are conducted using 10 GbE NICs unless otherwise noted. Watermarking and motion detection require additional vCPU and memory.

For all the tests, the virtual CPU (vCPU), memory, and network were configured according to Milestone best practices. The VMware vSphere configuration was in accordance with the VMware Compatibility Guide (www.vmware.com/resources/compatibility/search.php). In addition, Dell EMC PowerPath™ is recommended for use with block storage (FC and iSCSI) implementations on VNX and VNXe arrays, while Microsoft MPIO is recommended for use with Unity arrays.

The Dell EMC Surveillance Lab's host hardware met and exceeded the minimum system requirements for an ESXi/ESX installation. The Milestone Recorder VM was running on an ESXi 6.5 host using Dell EMC servers such as the Dell EMC FC640s, FC430s, R640s and Cisco UCS Blade Servers. 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.

VMware ESXi environment

We make sure our test environment host hardware meets the system requirements for an ESXi installation.

The VMware vSphere system requirements are as follows:

- Dell EMC PowerPath for block storage (FC and iSCSI)

For a list of compatible hardware, refer to the [VMware Compatibility Guide](#) and *Milestone XProtect Corporate System Requirements*.

Unity and SC series arrays

With the Milestone XProtect solution, a Unity or VNX family array is used for the Live DB, Archive DB, and VMware Datastores.

Unity or SC series can be deployed with GigE or 10 GbE iSCSI NICs. Dell EMC PowerPath™ multipathing is recommended for block storage like SC series. PowerPath multipathing automates data path management, failover and recovery, and optimized load balancing to ensure application availability and performance. Native operating system multiple path I/O (MPIO) is recommended for Unity arrays.

Isilon clustered storage system

The Archive DB can reside on Isilon scale-out network-attached storage (NAS). Isilon storage 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. Dell EMC recommends using an Isilon clustered storage system that is composed of five nodes or more.

Dell EMC ECS Object Storage

Dell EMC ECS is a complete software-defined cloud storage platform that supports the storage, manipulation, and analysis of video surveillance and unstructured data on a massive scale on commodity hardware. ECS is specifically designed to support the mobile, cloud, and Big Data workloads that are similar to large-scale workloads.

ECS provides UI, RESTful API, and CLI interfaces for provisioning, managing, and monitoring storage resources. Storage services provided by the unstructured storage engine (USE) ensure that video is available and protected against data corruption, hardware failures, and data center disasters. The USE enables global namespace management and replication across geographically dispersed data centers and enables the following storage services:

Object service

Enables you to store, access, and manipulate video and unstructured data. The object service is compatible with existing Amazon S3, Dell EMC Centera™ content addressable storage (CAS), and Atmos™ APIs.

Hadoop Distributed File System (HDFS)

Helps you use your ECS infrastructure as a Big Data repository against which you can run Hadoop analytic applications.

The provisioning service manages the provisioning of video surveillance storage resources and user access. Specifically, it handles user management, authorization, and authentication for all provisioning requests, resource management, and multitenancy.

You can scale up, scale out, and add users, applications, and services, as well as manage your local and distributed storage resources for your surveillance data through a single view.

Dell EMC CIFS-ECS

CIFS-ECS is a lightweight application that allows you to upload and download files to a Dell EMC ECS storage platform. It creates a Windows virtual drive to ECS cloud storage and transfers data from a Windows platform to an ECS using REST S3 API. CIFS-ECS is designed as an easy access to data in the cloud by allowing Windows applications to interface with an ECS storage server through standard file system APIs.

ECS combined with CIFS-ECS provides applications and users efficient access to content in the cloud from a Windows platform.

Dell Embedded Box PCs

Dell Embedded Box PCs are ruggedized, fanless, highly reliable devices for a variety of use cases, including process and discrete manufacturing, fleet management, kiosks, digital signage, surveillance and automated retail solutions. Embedded computers must run reliably 24x7 for extended deployments, and withstand higher and lower temperatures than ordinary PCs in environments that can bring high amounts of shock, vibration, moisture and high electromagnetic radiation.

Dell Embedded Box PCs are designed to MIL-STD-810G standards, building on the expertise of our rugged device engineers. Flexible, with many input/output (I/O) options, they run on powerful multicore Intel® processors. PCI/PCIe card slots provide flexible expansion for new I/O and graphics capabilities.

Dell EMC PowerEdge servers

Dell PowerEdge™ servers are ideal for recording and managing terabytes of video from distributed locations. This section describes best practices for configuring a PowerEdge server system for this solution.

PowerEdge single-unit rack servers are excellent servers when external NAS clusters or block arrays are planned for surveillance storage.

PowerEdge dual-unit rack servers are designed for local video storage where external surveillance storage will not be used.

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:

- [Sizing guidelines](#).....20
- [Volume limits](#).....20
- [Live DB only sizing](#)..... 20
- [Archive DB sizing](#)..... 21
- [Dell EMC ECS](#) 24
- [Dell Embedded Box PCs](#).....24
- [Dell EMC PowerEdge server storage](#)..... 25

Sizing guidelines

For best performance and usability, XProtect and each storage tier must be correctly configured. This section describes the results of the solution testing as well as configuration settings we found important during our lab tests.

In the same way that the configuration and performance of one production implementation varies from another, lab tests also vary among production environments. Use the performance statistics and configuration information presented here as the base guideline.

For information about the test objectives and the test procedure, see [Testing and validation overview](#).

Note

While all tests for this solution were performed on Milestone XProtect Corporate, the same codebase is used by XProtect Enterprise. Therefore, all sizing data in this document is valid for XProtect Enterprise also. XProtect Enterprise, however, limits the number of recording servers supported.

Dell EMC Surveillance Lab testing shows that Storage Pools provide better performance and are easier to manage for the end customer. Therefore, validation tests are only conducted using Storage Pools. Although all the underlying Storage Pool's RAID Groups configurations perform equally with the surveillance workload, we recommend RAID 6 (6+2) for resilience.

To maximize caching performance for surveillance workloads, the Dell EMC Surveillance Lab recommends using two SSD system drives per node in clusters where it is supported, such as the NL-series.

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.

Live DB only sizing

The Live DB can reside on DAS or block (FC or iSCSI) storage.

For Live DB sizing we recommend the following:

- Set the Live DB retention period from 2 to 24 hours, depending on the implementation's requirements.
- To allow space for video file collection and periodic archiving from the Live DB to the Archive DB, size the available Live DB storage to be at least twice the duration that video is configured to reside on the Live DB. However, the size of the Live DB can be set to a greater value to meet the needs of a particular implementation. The additional space is to accommodate network maintenance, traffic congestion, and other conditions that may temporarily restrict the bandwidth in an IP network and take time to resolve. For a conservative implementation, you may want to include additional space to allow for break-fix conditions.

Note

The write block size is determined by how the LUN or disk was formatted. For more information, see the Hard disk formatting section in *Dell EMC Storage with Milestone XProtect: Configuration Guide*.

The following table shows details of Live DB performance on Unity arrays during lab testing.

Table 2 Live DB Unity performance

Array	Milestone Version	Disks	RAID	No of Recorders	Recorder BW (MB/s)	Array BW (MB/s)		Maximum Capacity (RAW)
						iSCSI	FC	
Unity300 ^a	2018	80	6	7	42	294	323	2.34 PB
		150	6	13	42	546	601	
Unity400 ^a	2018	120	6	10	42	420	462	3.9 PB
		250	6	21	42	882	970	
Unity500	2018	100	6	9	42	378	416	7.8 PB
		200	6	18	42	756	832	
		400	6	36	42	1512	1663	
		500	6	36	42	1512	1663	
Unity600 ^a	2018	100	6	9	42	378	416	9.7 PB
		200	6	18	42	756	832	
		400	6	36	42	1512	1663	
		500	6	45	42	1890	2079	
		1000	6	45	42	1890	2079	

a. These values are extrapolated from Dell EMC Surveillance Lab test results.

Archive DB sizing

Starting with XProtect Corporate release 2013, you can use either file or block storage for the Archive DB. If your installation is an earlier version of XProtect, and you want the Archive DB to reside on an Isilon cluster running OneFS 7.0 (or later), check with Milestone to see if they provide a patch for your specific XProtect release.

When the Archive DB resides on EMC VNX block storage, the LUN or disk being formatted determines the write block size. For more information, see the Hard disk formatting section in *Dell EMC Storage with Milestone XProtect: Configuration Guide*.

The following tables provide information to help you understand the sizing components of NAS- and SAN-based implementations. Isilon sizing was based on five recording servers writing to a single Isilon node.

Table 3 Archive DB with Unity

Array	Milestone version	Disks per pool	RAID	BW per Recorder (MB/s)	No. of Recorders	Array BW (MB/s)		Maximum (RAW)
						iSCSI	FC	
Unity 300	2018	80	6	34	7	306	336	2.34 PB
		150	6	34	16	544	598	
Unity 400 ^a	2018	120	6	42	14	476	523	3.9 PB
		250	6	42	28	952	1047	
Unity 500 ^a	2018	100	6	42	11	374	411	7.8 PB
		200	6	42	22	748	882	
		400	6	42	45	1530	1683	
		500	6	42	45	1530	1683	
Unity 600 ^a	2018	100	6	42	11	374	411	9.7 PB
		200	6	42	22	748	822	
		400	6	42	45	1530	1683	
		500	6	42	56	1904	2094	
		1000	6	42	56	1904	2094	

a. Extrapolated based on Dell EMC Surveillance Lab test results.

Note

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

The following table shows details of Archive DB performance on SC series arrays during lab testing.

Array	Milestone Version	Disks	RAID	No of Recorders	Recorder BW (MB/s)	Array BW (MB/s) iSCSI	Maximum Array Capacity (RAW)
SC3000	2018	75	6	9	32	288	1 PB
		150	6	18	32	576	
		222	6	27	32	864	

Table 4 Archive DB: Isilon sizing

Array	Milestone version	OneFS	Recorders per node	Avg BW per node (MB/s)	Avg BW per recorder (MB/s)	Cluster size	Disk type
A200	2018	8.1.0 ^a	1:1	40	40	4	4 TB, +2:1
			2:1	60	30	4	4 TB, +2:1

Table 4 Archive DB: Isilon sizing (continued)

Array	Milestone version	OneFS	Recorders per node	Avg BW per node (MB/s)	Avg BW per recorder (MB/s)	Cluster size	Disk type
A2000	2018 #GUID-3638D3 3C-97F1-4ECE- 9854-02E8B62 710E1/FN- EXTRAPOLATE	8.1.0 ^a	1:1	33	33	4	10 TB, +2:1
			2:1	66	33	4	10 TB, +2:1
	2017	8.1.0 ^a	1:1	33	33	4	10 TB, +2:1
			2:1	66	33	4	10 TB, +2:1
NL410	2017	8.0.1	1:1	33	33	4	4 TB, +2:1
			2:1	66	33	4	4 TB, +2:1
			3:1	99	33	4	4 TB, +2:1
			4:1	132	33	4	4 TB, +2:1
HD400	2017 ^b	8.0.1	1:1	33	33	4	6 TB, +2:1
			2:1	66	33	4	6 TB, +2:1
			3:1	99	33	4	6 TB, +2:1
			4:1	132	33	4	6 TB, +2:1
	2014	7.2	1:1	37	37	5	6 TB, +2:1
NL400	2014	7.2	2:1	102	51	4	4 TB, +2:1
			4:1	164	41	4	4 TB, +2:1
		7.1	1:1	40	40	5	1 TB, +2:1
			3:1	120	40	5	1 TB, +2:1
			4:1	160	40	5	1 TB, +2:1
			5:1	194	38	5	1 TB, +2:1
X400	2014	7.1	1:1	34	34	5	1 TB, +2:1 ^c
			2:1	68	34	5	1 TB, +2:1
			4:1	136	34	5	1 TB, +2:1

- a. See Dell EMC Storage with Milestone XProtect Configuration Guide for additional information regarding OneFS 8.1 installations.
b. Extrapolated based on Dell EMC Surveillance Lab test results.
c. The drive size is not relevant for performance considerations.

Note

Add disk drives are SATA 7200 RPB unless otherwise noted.

The maximum bandwidth per LUN for the Live DB was limited to 34 MB/s (iSCSI) or 40 MB/s (FC)-see the previous table, Live DB VNX performance. Therefore, the maximum bandwidth between the recording servers and Isilon is based on those

values. Future tests will validate higher Archive DB bandwidth capabilities for greater numbers of recording servers per node.

Dell EMC ECS

The test results are based on a model in which the constant-bandwidth surveillance video traffic remained unaffected during select storage failure scenarios, such as disk rebuild, node failures, and failing network paths.

We performed all tests with disk drive failures, node failures, storage process failures, or NIC failures to ensure a worst-case scenario for all sizing parameters.

Dell EMC recommends:

- Using SSD or 15k rpm SAS drives for the CIFS-ECS cache disks.
- Calculating drive space requirements for local disk and ECS buckets based on the retention times used.

Note

For milestone recorders, the Archive DB disk must be configured as the cache disk for the CIFS-ECS tool.

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

Table 5 Dell EMC ECS Object Storage test results

Cluster	ECS version	No. of drives/ECS node	ECS node drives		Recorders per node	Bandwidth (MB/s)	
			Size	Type		Recorder	Node
U400	3.2.0.0	30	8 TB	NL-SAS	1	40	40
		30	8 TB	NL-SAS	2	40	80
		30	8 TB	NL-SAS	3	26	78
		30	8 TB	NL-SAS	4	20	80

Dell Embedded Box PCs

We conducted functional tests to determine how Milestone works with Dell Embedded PCs.

Table 6 Dell Embedded PC test results

Option	Dell Embedded Box PC 5000		
	Live DB only to Local disk	Live DB only to External storage	Live DB to local disk and Archive DB to External storage
Video Data storage	Local	External storage	Local -> External
Drive type and Size	2 TB, SATA Flash drive	ISILON - NAS	Local SSD -> ISILON NAS

Table 6 Dell Embedded PC test results (continued)

Option	Dell Embedded Box PC 5000		
	Live DB only to Local disk	Live DB only to External storage	Live DB to local disk and Archive DB to External storage
RAID Used	0	NA	
CPU	Intel Core i7 Processor	Intel Core i7 Processor	Intel Core i7 Processor
RAM GB	32	32	32
Network speed	2 X 1 Gb	2 X 1 Gb	2 X 1 Gb
XProtect Version	2018	2018	2018
Recorder write BW	10 MBps	32 MBps	10 MBps local DB - > 15 MBps Archive DB
OS	Windows 10 Enterprise 64Bit	Windows 10 Enterprise 64Bit	Windows 10 Enterprise 64Bit

Dell EMC PowerEdge server storage

The test results are based on a model in which the constant-bandwidth surveillance video traffic remained unaffected during select storage failure scenarios, such as disk rebuild, failing processors, and failing network paths.

We performed all tests with disk drive failures and NIC failures to ensure a worst-case scenario for all sizing parameters.

With local storage we recommend using RAID 6 to accommodate the disk rebuild duration. For example, in the Dell EMC Surveillance Lab, a disk rebuild in a 16 x 10 TB disk system takes at least 36 hours to complete when 10TB of data and 37.5 MB/s of write are in place.

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

Table 7 Dell EMC PowerEdge server test results

Server model	XProtect Version	Drive type	Drive size	RAID	Write BW (MB/s)	CPU (%)	IOPS	Rebuild duration	Data storage	
									Live DB (tier-1)	Archive DB (tier-2)
R740xd	2018	NL-SAS	16 x 4 TB	6	90	12	180	8 days	Local	n/a
		NL-SAS	16 x 10 TB	6	90	12	180	8 days	Local	
R730xd	2017	NL-SAS	8 x 4 TB	6	68	8	128	35 hrs	Local	n/a
		NL-SAS	15 x 4TB	6	137	18	117	33 hrs	Local	
		NL-SAS	15 x 10 TB	6	136	10	105	7.2 days	Local	
R630 *	2017	SSD	5 x 1 TB	5	69	10	142	1.5 hrs	Local	Isilon
R530	2017	NL-SAS	6 x 6 TB	6	50	60	65	20 hrs	Local	n/a

Table 7 Dell EMC PowerEdge server test results (continued)

Server model	XProtect Version	Drive type	Drive size	RAID	Write BW (MB/s)	CPU (%)	IOPS	Rebuild duration	Data storage	
									Live DB (tier-1)	Archive DB (tier-2)
R430	2017	SSD	2 x 400 GB	1	87	13	126	n/a	SCV2080	n/a
R430	2017	SSD	2 x 400 GB	1	86	13	126	n/a	Local	Isilon

* These values are extrapolated from Dell EMC Surveillance Lab test results.

Note

- All tests performed on local storage are Milestone certified.
 - Stripe size of 128 was used for video storage virtual disks.
 - Servers used for external storage must meet Milestone's minimum requirements.
-

Note

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

CHAPTER 5

Testing and validation

This chapter describes the testing used to validate this solution.

- [Test objectives](#).....28
- [Test parameters](#).....28
- [Tests conducted](#)..... 28

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

Test parameters

All test parameters and scenarios reflect standard production behavior for Milestone XProtect Corporate 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 Milestone 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, video is simultaneously recalled or reviewed at a rate equal to 20 percent of the write rate. Tests are run with the SmartConnect™ configuration in place and the SMB shares are 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 Unity and SC series block storage arrays were protected using RAID with hot spare disks. For the test, disk failure scenarios were induced and the data rebuild to the hot spare disks was observed with effect to write bandwidth.

The Isilon cluster was protected using a +2 protection scheme that allows for two simultaneous disk failures. For the test, two disks are failed and then recovered. The

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

Disk failure test with ECS

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.

ECS employs a hybrid model of triple mirroring data, metadata, and indexing. Erasure coding is also used for enhanced data protection and reduction of storage overhead. For data integrity, ECS uses checksums.

When the system labels a drive as FAILED, the data protection logic rebuilds the data on that drive on other drives in the system. The FAILED drive no longer participates in the system in any way. ECS requires a minimum of four nodes to be able to conduct the default erasure coding and six nodes for the cold archive option.

The disk rebuild operation did not affect the write rate, video quality, or result in dropped video.

The Unity and SC series block storage arrays were protected using RAID with hot spare disks. For the test, disk failure scenarios were induced and the data rebuild to the hot spare disks was observed with effect to write bandwidth.

The Isilon cluster was protected using a +2 protection scheme that allows for two simultaneous disk failures. For the test, two disks are failed and then recovered. The SmartFail process started and the CPU utilization of the node increased with no observed effect to the write streams.

NIC failure test

The Unity and SC series block storage arrays were configured with multiple paths to the recorders using Microsoft MPIO. Multiple NICs were configured with the recorders and controllers for redundancy. The Unity and SC series hard NIC failure test removes one nic cable from the array. Recorders that were configured with multipathing reconnected to the volume across another available path. To reduce the reconnection time and eliminate video loss, adjust the TCP retransmission timers. For more information, see the [Dell EMC Block Storage with Video Management Systems Configuration Best Practices Guide](#).

The Isilon hard NIC failure test removes 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 [89232](#), *Configuring sysctls and making sysctl changes persist through node and cluster reboots and upgrades* 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 ECS

The ECS hard NIC failure test removes one NIC cable from the active node that was involved in active recording to simulate the NIC failure scenario.

The Dell EMC Surveillance Lab uses two 10 GbE, 24-port or 52-port Arista switches that are used to transfer data to and from customer applications as well as internal node-to-node communications. These switches are connected to the ECS nodes in the same rack and employ the Multi-Chassis Link Aggregation (MLAG) feature, which logically links the switches enabling active-active paths between the nodes and customer applications. This configuration results in higher bandwidth while preserving resiliency and redundancy in the data path. Any networking device supporting static LAG or IEEE 802.3ad LACP can connect to this MLAG switch pair. Because the switches are configured as MLAG, these two switches appear and act as one large switch.

The NIC failure tests did not affect the write rate, video quality, or result in dropped video.

NIC Failure test with NIC aggregation in Active/Passive

The hard NIC failure test with Active/Passive aggregation was run 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.

TCP transmission timers can be adjusted to reduce the reconnection times during Nic failures on recorders that use Microsoft MPIO. For more information, see the [Dell EMC Block Storage with Video Management Systems Configuration Best Practices Guide](#).

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

An unexpected single node hard failure was simulated, which causes the servers that were writing to that node to reconnect to a new node.

During the 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.

Also, 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
```

Note

During an abrupt failure of a node, the recorders writing to that node reconnect to SmartConnect and can buffer the video during reconnection. Data that was already written to the disk cannot be recovered which can range from about 4 to 5 seconds.

WARNING

If running a mixed workload, these changes can adversely affect the other workloads that might be present on the cluster.

Node poweroff test with ECS

ECS employs a hybrid model triple mirroring data, metadata, and indexing. Erasure coding is also used for enhanced data protection and reduction of storage overhead.

Erasure coding provides enhanced data protection from a disk or node failure that is storage efficient as compared to conventional protection schemes. The ECS storage engine implements the Reed Solomon 12+4 erasure-coding scheme, in which a chunk is broken into 12 data fragments and 4 coding fragments for parity. These 16 fragments are then dispersed across nodes at the local site. The data and coding fragments for each chunk are equally distributed across nodes in the cluster. For example, with 8 nodes, each node stores 2 of the 16 fragments. The storage engine can then reconstruct a chunk from any 12 fragments of the original 16.

One of the ECS nodes was manually shutdown. The CIFS-ECS tool load balanced the traffic across all the available nodes and the recorders bypassed the failed node. The node failure did not affect the write rate, video quality, or result in dropped video.



If running a mixed workload, these changes can adversely affect the other workloads that might be present on the cluster.

Node reboot test with ECS

One of the ECS nodes was manually restarted to simulate a node reboot. The CIFS-ECS tool load balanced the traffic across all the available nodes and the recorders bypassed the failed node. The node reboot did not affect the write rate, video quality, or result in dropped video.

CHAPTER 6

Conclusion

This chapter summarizes the testing for this solution:

- [Summary](#)34

Summary

The Dell EMC Surveillance Lab performed comprehensive testing with Milestone XProtect Corporate to benchmark application performance in a Milestone tiered storage environment. Milestone engineering has confirmed that the test results achieved are consistent with the GA versions released.

Testing focused on the Archive DB enhancements of XProtect 2018 that allow the use of Isilon clusters running OneFS 7.0 or later. The archive process, when writing to an Isilon cluster, easily handled all video accumulated between archive process executions. In addition, various forced Isilon failures did not affect the Live DB to Archive DB process.

When using an Isilon cluster for the Archive DB, only two modifications to the XProtect configuration are necessary: increase the number of archive process threads and increase the write block size to the Archive DB.

We also tested an FC- and iSCSI-attached VNX for the both Live DB and Archive DB. The results for this configuration represent the maximum tested, not the array maximum.

For both FC- and iSCSI-attached storage, formatting with 8192 KB blocks for Live DB and 64 KB blocks for Archive DB is required.

Testing and validation of this solution produced these key findings:

- FC- or iSCSI-attached VNX can be used in place of internal server storage for both Live DB and Archive DB video storage.
- When connectivity to an Isilon cluster node fails, the XProtect Corporate Recorders attached to that node can experience some data loss. The Dell EMC Surveillance Lab observed data loss of approximately 15 seconds as the recorder failed to another node.
- Isilon clusters can be used as Archive DB video storage starting with XProtect Corporate 2016.
- One SMB share per XProtect Corporate recording server is required.
- Bandwidth to the Isilon clusters and VNX arrays for the Archive DB was not affected during numerous forced failures on the cluster.

Dell EMC Unity and SC series arrays

The use of storage pools to create LUNs within the Dell EMC Unity or SC series 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 Recorder view of the storage is based on the assigned quota and not the entire file system. We recommend using SmartQuotas with Milestone XProtect Corporate as a best practice.