

Enabling IP video surveillance solutions using Milestone XProtect Corporate, IBM Spectrum Scale and IBM Elastic Storage Server

A technical report

Overview

Challenge

How do you configure and size the storage system for extremely scalable IP video surveillance solutions using Milestone XProtect Corporate?

Solution

This paper provides various storage configuration guidelines to gain the best performance from IBM Spectrum Scale and Elastic Storage Server for larger camera video environments.

The objective of this paper is to discuss the storage considerations to assist users in properly sizing to operate in a very large digital video surveillance solution. The paper covers the test strategy used and the results of the testing. This paper describes a highly scalable architecture solution for deploying larger Milestone security solutions on IBM Spectrum Scale and Elastic Storage server (ESS). This paper provides recommendations, and configuration and tuning adjustments to help ensure an efficient installation of the joint solution with acceptable performance that meets the requirements of the organization.

The Milestone XProtect video management software is a powerful surveillance platform that is easy to manage. The configuration capabilities of the software design enable you to customize the solution to fit your requirements.

IBM® system engineers thoroughly tested the solution using the system parameters outlined in this paper. This technical white paper covers the test results that describe the performance of Milestone XProtect Corporate running on IBM Spectrum Scale™ and IBM Elastic Storage™ Server. It illustrates the hardware impact (on the processor, storage, and network) when changing the individual parameters of the video surveillance system, and provides recommendations on how to configure the surveillance system for optimal performance.



Assumptions and prerequisites

This solution paper assumes familiarity with the following technology skills:

- Proficiency in installing and configuring the Milestone XProtect Corporate suite of products.
- Familiarity with installing and configuring the IBM Spectrum Scale and Elastic Storage Server systems.
- Proficiency in installing and configuring Microsoft® Windows® and Linux® operating systems.

The configurations documented in this paper are based on the tests that were conducted in the IBM internal lab. These results might differ from individual production implementations in the field.

IBM Spectrum Scale

IBM Spectrum Scale was formerly known as IBM General Parallel File System (IBM GPFS™).

A high-performance enterprise platform for optimizing data and file management, Spectrum Scale is used extensively across industries worldwide. Spectrum Scale simplifies data management with integrated tools designed to help organizations manage petabytes (PB) of data and billions of files, and also control the cost of managing these ever-growing data volumes.

Part of the IBM Spectrum Storage® family, Spectrum Scale introduces revolutionary new features that clearly demonstrate the commitment of IBM to provide groundbreaking storage solutions, including:

- File encryption and secure erase
- Transparent flash cache
- Network performance monitoring
- Active file management (AFM) parallel data transfers
- Network File System (NFS) version 4 support and data migration
- Backup and restore improvements
- File Placement Optimizer (FPO) enhancements

IBM Spectrum Scale allows a group of computers to concurrently access a common set of file data over a direct attached or shared storage area network (SAN) infrastructure, a TCP/IP network, or a combination of connection types. The computers can run a mix of IBM AIX®, Linux or Microsoft Windows operating systems. Spectrum Scale provides storage management, information lifecycle management tools, centralized administration, and

shared access to file systems from remote Spectrum Storage clusters providing a global namespace.

A Spectrum Scale cluster can be a single node in a tiered storage solution, two nodes providing a high-availability platform supporting a database application, or thousands of nodes used for applications such as weather-pattern modeling. A Spectrum Scale file system is built from a collection of storage devices that contain the file system data and metadata. A file system can be built from a single disk or contain thousands of disks storing petabytes of data. Each file system can be accessible from all the nodes within the cluster. There is no practical limit on the size of a file system. The architectural limit for a single file system is more than a yottabyte. Some Spectrum Scale customers use single file systems up to 18 PB in size, while others use file systems containing billions of files. Applications access files through standard Portable Operating System Interface (POSIX) file system interfaces. Because all nodes see all file data, any node in the cluster can concurrently read or update a common set of files and thus enabling applications to scale out easily. Spectrum Scale maintains the coherency and consistency of the file system using sophisticated byte-range locking, token (distributed lock) management, and journaling. This approach means that applications using standard POSIX locking semantics do not need to be modified to run successfully on Spectrum Scale.

Spectrum Scale provides unparalleled I/O performance for unstructured data by:

- Striping data across multiple disks attached to multiple nodes
- Employing high-performance metadata (inode) scans
- Supporting a wide range of file system block sizes to match I/O requirements
- Using advanced algorithms to improve read-ahead and write-behind I/O operations
- Using block-level locking, based on a sophisticated scalable token management system, to provide data consistency while allowing multiple application nodes concurrent access to files

When creating a Spectrum Scale file system, raw storage devices are assigned to the file system as Network Shared Disks (NSDs). After an NSD is defined, all the nodes in the Spectrum Scale cluster can access the disk using a local disk connection or the NSD network protocol for shipping data over a TCP/IP or InfiniBand® connection.

Spectrum Scale token (distributed lock) management coordinates access to NSDs, helping to ensure the consistency of file system data and metadata when different nodes access the same file. Token management responsibility is dynamically allocated among designated manager nodes in the cluster. Spectrum Scale can assign one or more nodes to act as token managers for a single file system, allowing greater scalability for large numbers of files with high transaction workloads. In the event of a node failure, token management responsibility is transparently moved to another node.

All data stored in a Spectrum Scale file system is striped across all storage devices within a storage pool (irrespective of whether the pool contains 2 or 2,000 storage devices). When storage devices are added to a storage pool, existing file data can be redistributed across the new storage to improve performance. Data redistribution can be scheduled or can be done organically when there is a high data change rate. When redistributing data, a single node can be assigned to perform the task (to control the impact on a production workload). Alternately, all nodes in the cluster can participate in data movement (in order to complete the operation as quickly as possible).

Along with distributed token management, Spectrum Scale provides scalable metadata management by allowing all nodes of the cluster accessing the file system to perform file metadata operations. This feature distinguishes Spectrum Scale from other cluster file systems, which typically have a centralized metadata server handling fixed regions of the file namespace. A centralized metadata server can often become a performance bottleneck for metadata-intensive operations, limiting scalability and possibly introducing a single point of failure. Spectrum Scale solves this problem by enabling all nodes to manage metadata.

Spectrum Scale supports POSIX and NFS v4 access control lists (ACLs). NFS v4 ACLs can be used to serve files using NFS v4, but can also be used in other deployments (for example, to provide ACL support to nodes running Windows). To provide concurrent access from multiple operating system types, Spectrum Scale allows you to run mixed POSIX and NFS v4 permissions in a single file system and map user and group IDs between Windows and Linux/UNIX® environments.

IBM Elastic Storage Server (ESS)

Elastic Storage Server at glance

- IBM Power S822L storage servers
 - IBM Spectrum Scale for cluster management.
 - Models
 - GS1, GS2, GS4, GS8
 - GL2, GL4, GL6
 - Maximum of 348 NL-SAS disk drives
 - Maximum Usable capacity is 2PB in 42U rack space.
 - RAID- 1, RAID- 6 (8+2P or 8+3P) with IBM declustered RAID technology.
 - Scale capacity and/or performance in modular building blocks with automatic load balancing
 - Faster rebuild failed disks with IBM-developed erasure-coding based declustered RAID technology.
 - Single global name space for the file access.
-

IBM Elastic Storage Server is a modern implementation of software-defined storage (SDS), combining IBM Spectrum Scale software with IBM POWER8® processor-based servers and storage enclosures.

IBM Spectrum Scale, formerly known as IBM GPFS, is a parallel file system that is at the heart of IBM Elastic Storage Server. IBM Spectrum Scale scales system throughput with each new server while still providing a single namespace to the clients. This ability eliminates data silos and simplifies storage management. By consolidating storage requirements across your organization onto IBM Elastic Storage Server, you can reduce inefficiency and acquisition costs while simplifying management and improving data protection.

The key capabilities of IBM Elastic Storage Server include:

Software RAID

IBM Spectrum Scale Redundant Array of Independent Disks (RAID) runs IBM disks in a dual-ported storage enclosure that does not require external RAID storage controllers or other custom hardware RAID acceleration.

Declustering

IBM Spectrum Scale RAID distributes client data, redundancy information, and spare space uniformly across all disks of just a bunch of disks (JBOD). This distribution reduces the rebuild or disk failure recovery process penalties compared to conventional RAID. Critical rebuilds of failed multi-terabyte drives full of data can be accomplished in minutes, instead of hours or even days when using legacy RAID technology).

Data redundancy

IBM Spectrum Scale RAID supports highly reliable two-fault-tolerant and three-fault-tolerant Reed-Solomon parity codes (erasure coding) as well as three-way and four-way replication.

Large cache

Using a combination of internal and external flash devices along with large memory cache in the IBM Power® server, IBM Elastic Storage Server is better able to mask the inefficiencies and long latency times of nearline serial-attached SCSI (SAS) drives, while still using the high density of the drives themselves.

Superior streaming performance

The system can deliver over 25 GBps of sustained performance.

Scalability

As input/output (I/O) server configurations are added to an installed configuration, both capacity and performance scale while still providing a single namespace to clients. This means that installations can start small, and grow as data needs to expand.

Key components

This section provides the details about the components used in the solution.

Key lab components

- Milestone XProtect Corporate 2016
 - Axis Camera simulator
 - Windows 2012 Server
 - VMware ESX hypervisor
 - 10 Gb Ethernet Switch
 - IBM Spectrum Scale 4.2.1
 - IBM Elastic Storage Server GL6
 - Mellanox Ethernet adapters and switches
-

Milestone XProtect Corporate

Video management system (VMS) is a key component of a video surveillance solution that collects video data from the cameras and other data sources. Milestone XProtect Corporate is a leading open platform software designed to meet the highest security demands of large enterprise customers.

- **XProtect Corporate management server** is the central component of the VMS and is responsible for handling the system configuration, distributing configuration to other system components (such as recording servers), and for facilitating user authentication. The configuration data is stored in a standard Microsoft SQL server installed either on the management server itself or on a separate dedicated server. Management server provides an ability to manage geographically dispersed sites easily from a central location.
- **XProtect recording server** records and stores video, audio, and metadata, and provides operators and other users access to live and recorded streams. It receives the incoming camera data and stores the video content to the storage system.
- **XProtect Smart Wall** is an advanced video wall solution included with the software. It enables control room operators to share a wide range of information, including live and recorded video, images, maps, alarms, and text messages to better coordinate response activities.
- **XProtect Smart Client** is the primary user interface for security operators and other daily users. Its adaptable user interface provides an extremely efficient working environment that can be optimized for different tasks and operator requirements.

Server platform

Server platform plays a key role in video surveillance architecture. When designing and deploying large-scale VMS systems, Milestone recommends that you base the server design and implementation on virtualization technologies such as VMware vSphere. Virtualization technology offers a range of benefits when used in standard IT installations. Hypervisors such as VMware helps to host multiple servers on one physical server depending on the processor and memory resources.

Storage system

Storage system is the critical component in the surveillance system for storing the video contents. Ease of use, performance, scalability, and affordability are some of the important considerations for selecting the video storage unit. IBM

Spectrum Scale and Elastic Storage Server provides cost effective and highly scalable platform to store and playback the video archives from the large-scale VMS systems configured with thousands of cameras. IBM Spectrum Storage family provides a modular architecture, allowing you to scale in capacity and performance bandwidth as per the future requirements.

Networking

Networking infrastructure plays an important role in connecting the incoming camera data ingest and storing the archived content to the IBM Spectrum Scale storage systems. 10 Gb Ethernet connectivity is recommended between camera network, archiver servers, and IBM Spectrum Scale storage systems.

Lab validation configuration

This section lists the hardware used in the lab for validating the functionality and sizing information for the IP video surveillance solution using the IBM Storwize system.

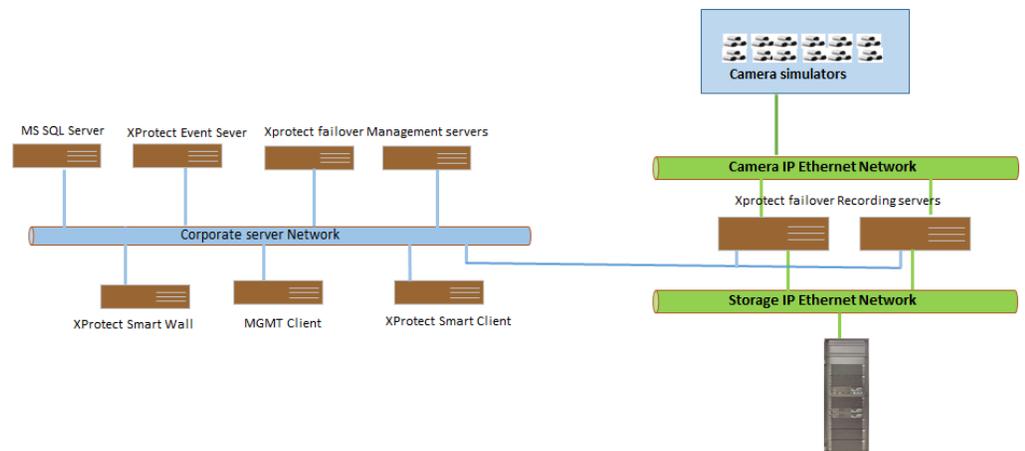


Figure 1: Milestone architecture with IBM Elastic Storage Server

No.	Role	Configuration
1	Milestone XProtect management server	Two processor cores, 8 GB memory, Windows 2012 R2, Milestone XProtect 2016 Management suite
2	Milestone XProtect recording servers	Four to six processor cores, 32 GB memory, two 10 Gb Ethernet, Windows 2012 R2, XProtect Corporate recording server, Cygwin package for 64-bit Windows, Spectrum Scale client for Windows platform
3	Camera simulator	One processor core, 8 GB memory, 10 Gb Ethernet, Windows 2012 R2 Axis Camera simulators were used to simulate the camera data. Each simulator server is configured with 150 cameras
4	Storage unit	IBM Elastic Storage Server GL6, three-way replicated metadata volumes, RAID 6 (8+2P) data volumes, layer 3+4 Ethernet network bonding

Table 1: Lab configuration details

10 Gb Ethernet network was used for ingesting the video data from camera network to the IBM Elastic Storage Server.

Testing and lab validation

Various tests were conducted in the lab with the focus on the storage related factors for the video archiving. The process used for testing the environment controlled and changed camera settings, and quantity of simulated cameras while monitoring key performance metrics for write latency, processor consumption, network utilization, and video frame loss. Axis camera simulator utility was used for simulating the cameras. Sustained continuous testing was performed with a given set of parameters for a given period that spans across multiple archive periods. Tests were also performed after the system had run for a period of four days with the initial benchmark value and some amount of data had accumulated in the archive database to simulate a production environment.

Test procedure

This section outlines the steps used for validating the solution in the lab.

IBM Elastic Storage Server configuration

This section lists the process to configure the IBM Elastic Storage Server.

1. **Create a cluster:** Run the `gssgencluster` command on the Elastic Management Server node to create the cluster. This command creates a Spectrum Scale cluster using all the nodes. The command assigns server licenses to each I/O server node. It applies the best practices of Spectrum Scale RAID configuration parameters. At the end of cluster creation, the SAS host bus adapter (HBA) firmware, storage enclosure firmware and drive firmware are upgraded if needed. This command can take a long time to run.
2. **Verify that the cluster is active:** Run the `mmgetstate` command from the I/O node to verify that the cluster is active and operational. Use the `-a` option to include all the nodes of the cluster.

```
[root@heliumems-hs ~]# mmgetstate -a
```

Node number	Node name	GPFS state
1	heliumio1-hs	active
2	heliumio2-hs	active
3	heliumems-hs	active
4	heliumpn1-hs	active
5	heliumpn2-hs	active
6	heliumpn3-hs	active
7	ISV512-Miles1	active

Figure 2: Spectrum Scale cluster status

3. **Create the recovery groups:** The `gssgenclusterrgs` command creates the recovery groups and the associated configurations.
4. **Verify the recovery configuration:** Log in to one of I/O group nodes and run the `mmfsrecoverygroup` to check the recovery group configuration.

```
[root@heliumems-hs ~]# mmfsrecoverygroup
```

recovery group	declustered arrays with vdisks	vdisks	servers
rg_heliumio1-hs	3	31	heliumio2-hs.tuc.stglabs.ibm.com, heliumio1-hs.tuc.stglabs.ibm.com
rg_heliumio2-hs	3	31	heliumio1-hs.tuc.stglabs.ibm.com, heliumio2-hs.tuc.stglabs.ibm.com

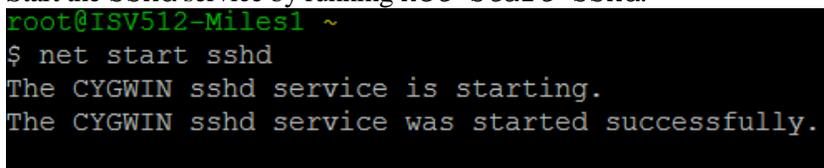
Figure 3: IBM Elastic Storage Server recovery group configuration

5. Run the following command:
`gssgengpfsvdisk -v 1 -b <block size for data> -m 1 -n <block size for metadata> -p 5 -l <Stanza file name>`
6. Edit the extent sizes in the stanza file created in the above step as per your requirement.
7. Create the VDisks by running the following command:
`mmcrvdisk -F <specify the stanza file name>`
8. Verify you have the data and metadata VDisks are created in each recovery group.
9. Create the NSD by running the following command:
`mmcrnsd -F <specify the same generated file name>`
10. Create the file system by running the following command:
`mmcrfs <file system name> -F <Stanza file name> -B <block size> --metadata-block-size <meta data block size> -j scatter -T <mount point>`
11. Mount the file system by running the following command:
`mmm mount <file system name> -a`
12. Check if the file system is mounted using `mmlsmount all`.

Spectrum Scale client configuration on Windows archiver server

This section describes the steps involved in installing and configuring the Spectrum Scale client on the Windows archive server.

1. Cygwin is a POSIX environment available for Windows and can be downloaded from the [Cygwin website](#). Download and install the Cygwin 64-bit Windows binaries on the Windows archiver server. Refer to IBM Knowledge Center for the detailed instructions on installing the Cygwin package.
http://www.ibm.com/support/knowledgecenter/STXKQY_4.2.1/com.ibm.spectrum.scale.v4r21.doc/bl1ins_windowscygwin.htm#windowscygwin
2. Install and configure the `sshd` service on the Windows Cygwin server using the `ssh-host-config` command.
3. Start the `sshd` service by running `net start sshd`.



```
root@ISV512-Miles1 ~  
$ net start sshd  
The CYGWIN sshd service is starting.  
The CYGWIN sshd service was started successfully.
```

Figure 4: Starting SSH service on the Cygwin server

4. Create a `root` user on the Windows server and add it to the local administrator group.
5. Configure the key-based password-less SSH connection between the Windows server and the Spectrum Scale storage cluster nodes. Make sure

- that the SSH connection is working without the password between the Windows client and the Spectrum Scale nodes in both directions.
6. Add the Windows archiver client to the Spectrum Scale cluster using the `mmaddnode` command. Run the command: `mmaddnode -N <windows client hostname>` command in the ESS I/O node.
 7. Configure the Spectrum Scale client license on the Windows server using the `mmchlicense` command. Run the command: `mmchlicense client -accept -N <Windows client hostname>`.
 8. Configure the Windows mount point on the Spectrum Scale file system using the `mmchfs` command. Run the command: `mmchfs <file system name> -t <Drive Letter>`.
 9. Start the Spectrum Scale service on the Windows node using the `mmstartup -N <node name>` command.
 10. Mount the Spectrum Scale storage file system on the Windows node using the `mmmount <file system name> -N <windows node name>` command. This will mount the Spectrum Scale file system as the drive letter on the Windows archiver server.

Milestone XProtect configuration

This section describes the steps required on the Milestone system.

1. Configure two servers for simulating camera data ingestion, and make sure that each server is configured with 150 cameras for data simulation.
2. Configure the `root` user as the log on account for the Milestone recording server service. This is required to grant the write permissions to the IBM Elastic Storage Server devices.

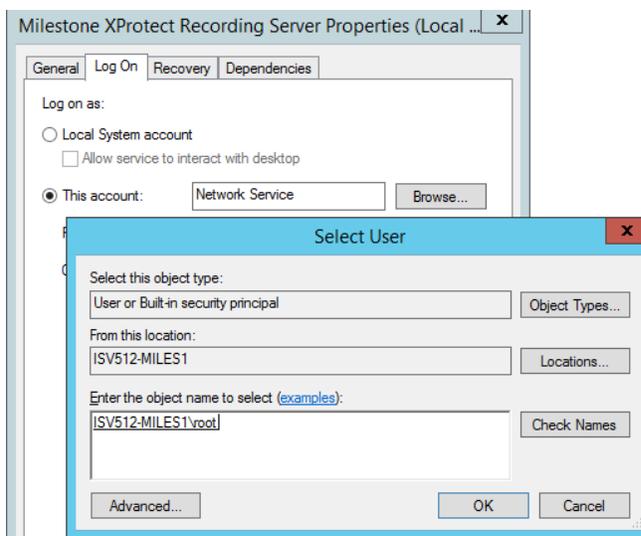


Figure 5: configuring the Milestone service as the root user

3. Configure live DB and archive storage locations to consume data from the

IBM Elastic Storage Server. Select the drive letters from the IBM Elastic Storage Server as the recording storage.



Figure 6: Milestone Storage configuration details

4. Configure the 150 cameras with HD 1080p, 20 frames per second on the recording server for the simulation.

Test strategy for validation

Axis camera simulator was used to simulate camera network. The Axis simulator camera count was set at 300 and the test was run for 5 days, without any loss of data frame.

Test settings

In the management client, select the Axis camera node under the recording server and click the Settings tab to make the following changes:

- Set **Codec** to **H264**
- Set **Compression** to **30**
- Set **Motion duration** to **0**
- Set **Frames per second** to **20**
- Set **Motion mode** to **Always**

After performing the test run for a couple of hours, the test team logged in to XProtect Smart Client from both XProtect Smart Client systems and performed the following steps.

1. Open a 24-camera view on each system and click the Playback tab.
2. Start the video loop playback.
 - a. Use the timeline selector to select the last hour of video.
 - b. Increase the playback speed by eight times.
 - c. Click Play.

During the 5-day test, 150 cameras were simulated with h.264, 1080p HD resolution, 20 frames per second (FPS), without any data overflow errors.

The following hardware fault injections were simulated at the storage level to measure the sustainability of the IBM storage system.

- Disk failures and rebuilds at the IBM Elastic Storage Server
- Rebooting of one of the ESS storage I/O node
- Network Interface Card (NIC) link failures for checking the active/active network bonding on the ESS

No media overflow errors were noticed during these extended testing.

Test results

Lab test results shows that 300, 1080p cameras produced very minimal workload on the Elastic Storage Server. However, as a result of this proof of concept testing with 300 cameras, IBM Spectrum Scale and Elastic Storage System is capable of providing the video storage for thousands of cameras. It is recommended to start with three to four recording servers, each configured with the 300-camera count as the base building block configuration and monitor the parameters before adding the additional recording servers and cameras, without the risk of losing data.

Two separate 10 Gb networks were configured on the recording server. One for the incoming camera ingestion and the second one for storing the data on the IBM Elastic Storage Server. During the performance run, the 10 Gb adapter operated at a very low utilization. In a moderate camera count environment, binding multiple 1 Gb adapters might be effective, but this was not tested.

Configuration guidelines

When choosing the hardware for a traditional digital video surveillance solution, the following criteria need to be evaluated:

- Number of cameras per server
- Frame rate and resolution of the cameras
- Retention period
- Physical location of the servers, cameras, and network infrastructure
- Bandwidth capabilities of the network
- Remote cameras or wide area networking for remote locations

The team conducted various tests in the lab to determine the sustainability of the IBM storage system with Milestone software.

Recording server

Refer to the following high-level guidelines for the recording server.

- It is recommended to configure maximum of 300 cameras per recording server.
- Four to six processor cores and 32 GB memory is recommended on the recording server.
- Two separate 10 Gb Ethernet interfaces: One for the camera video ingesting and the second one for the IBM Elastic Storage connectivity purpose is recommended
- A separate VLAN is recommended for the IBM Elastic Storage Server network connectivity to minimize the broadcast traffic.
- The following tuning parameters are recommended on the Milestone recording server:
 - `hardwarestartup maxDelay=1000`
 - `low_priority_archive_thread_pool_size =2`

IBM Elastic Storage Server

IBM Elastic Storage Server provides high performance and extremely scalable solutions suitable for the large video surveillance environments.

Elastic Storage Server model	GL6	GL4	GL2
Connectivity	10 Gb, 40 Gb, and InfiniBand	10Gb, 40 Gb, and InfiniBand	10 Gb, 40Gb, and InfiniBand
Maximum storage enclosures	6	4	2
Maximum drives	348	232	116

Figure 7: IBM Elastic Storage Server models summary

The following high-level recommendations are provided for Elastic Storage Server:

- IBM Elastic Storage Server GL family supports 2, 4, 6, and 8 TB NL-SAS drives and provides up to 2 PB of usable storage capacity in each building block.
- Four 10 Gb Ethernet ports are recommended on each IBM Storage Server I/O node.
- Active/Active Ethernet network bonding is recommended.
- 40 Gb and InfiniBand networks are recommended for very large environments.

- Three-way replicated VDisks for metadata and 8+2P RAID 6 VDisks for the data is recommended.
- In case of distributed islands of cameras, Storwize V5000 can be configured as the edge video storage system at the remote locations for storing the minimal data. At the core central location, larger scalable solutions based on IBM Spectrum scale and Elastic Storage Server can be configured to store the huge video content
- In the distributed model, network bandwidth becomes a significant factor and data placement and pruning policies might be needed to transfer the data effectively between various remote edge locations to the central core locations.

IBM Spectrum Scale configuration:

- It is recommended to create the Spectrum Scale GPFS file system with 4 MB or larger block size for the data and 256 KB for meta data purpose.
- Separate VDisks for metadata is recommended.
- Separate file sets are recommended for each archiving server for proper data classification.
- Larger page pool is recommended on the IBM Elastic Storage server I/O nodes.
- Minimum 4 GB of page pool size is recommended on the Windows archiver server.

Summary

IBM has performed extensive testing with Milestone XProtect Corporate and IBM Spectrum Scale and IBM Elastic Storage system. Depending on your requirements, Milestone and IBM storage system architecture allows a highly scalable and modular solution to hold the video data for the thousands of the camera. It also provides larger capacity solutions to hold multiple PB of storage required for the longer archiving purpose.

Software-defined storage: IBM Elastic Storage Server is a modern implementation of software-defined storage, combining IBM Spectrum Scale software with IBM POWER8 processor-based servers and storage enclosures. IBM Spectrum Scale scales system throughput with each new server while still providing a single namespace to the clients.

Designed for scalable and performance: Milestone XProtect Corporate, IBM Spectrum Scale, and IBM Elastic Storage system provides *best-in-class* and extreme scalable video solutions required for very large enterprises. This solution scales into PB of data with sustained streaming performance required for larger video surveillance environments.

Resources

The following links provide useful references to supplement the information contained in this paper:

- IBM Systems on PartnerWorld
ibm.com/partnerworld/systems
- IBM Redbooks
ibm.com/redbooks
- IBM Publications Center
ibm.com/e-business/linkweb/publications/servlet/pbi.wss?CTY=US

About the authors

Dave Taylor is an executive architect in IBM Systems Group ISV Enablement Organization. He has more than 25 years of experience working with the IBM System Storage and server platforms. You can reach Dave at dltaylor@us.ibm.com.

Sanjay Sudam is a consultant and an architect in IBM Systems Group ISV Enablement Organization. He has more than 15 years of experience working with the IBM System Storage and server platforms. You can reach Sanjay at Sanjay.Sudam@in.ibm.com.



© Copyright IBM Corporation 2016
IBM Systems
3039 Cornwallis Road
RTP, NC 27709

Produced in the United States of America

All Rights Reserved

IBM, the IBM logo and ibm.com are trademarks or registered trademarks of the International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked items are marked on their first occurrence in the information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

INFINIBAND, InfiniBand Trade Association and the INFINIBAND design marks are trademarks and/or service marks of the INFINIBAND Trade Association.

Other product, company or service names may be trademarks or service marks of others.

References in the publication to IBM products or services do not imply that IBM intends to make them available in all countries in the IBM operates.



Please recycle
