



Technical Report

Video Surveillance Solutions with NetApp E-Series Storage Sizing Considerations

Joel W. King, James Laing, NetApp
July 2013 | TR-4199

Abstract

Video surveillance solutions using E-Series storage offer the physical-security integrator a highly scalable repository for video management systems supporting high camera counts, megapixel resolutions, high frame rates, and long retention periods. The architecture is designed to provide high reliability and availability to meet the demands of video surveillance deployments.

TABLE OF CONTENTS

1	Introduction	4
1.1	Publication Scope	4
1.2	Audience	4
1.3	Training Offerings	4
2	Sizing Considerations	4
2.1	System Requirements	5
2.2	General Considerations	6
2.3	Retention Period	6
2.4	Reserve Capacity	7
2.5	Cameras	7
2.6	Centrally Stored Video Clips	14
2.7	Tiered Storage	14
2.8	High Availability	16
3	Sizing Fundamentals	17
3.1	Storage, Operating System, and File System Capacity Considerations	18
4	Video Surveillance Sizing Examples	20
4.1	Sizing Example 1: A Simple Deployment	20
4.2	Sizing Example 2: Larger System with Failover and RAID 10	22
4.3	Sizing Example 3: Complex Deployment for a Multiuse Center	24
5	Sizing Checklist	26
6	Summary	27
	Definitions	27
	References	28
	Version History	29

LIST OF TABLES

Table 1)	Training offerings	4
Table 2)	Usable capacity by RAID level	19
Table 3)	E-Series disk shelves for video surveillance deployments	19
Table 4)	Multiuse project	24
Table 5)	Data rate and storage per camera	25
Table 6)	Camera assignment per server	25

Table 7) Sizing solution26

LIST OF FIGURES

Figure 1) System requirements5
Figure 2) Example of throughput versus retention.....7
Figure 3) Axis Design Tool8
Figure 4) Video stream settings.....9
Figure 5) Daylight versus nighttime data rate11
Figure 6) OnSSI Ocularis storage configuration12
Figure 7) 64-camera transition from 1x to 16x.....14
Figure 8) OnSSI Ocularis ES recording and archiving configuration.15
Figure 9) Sizing fundamentals.....18
Figure 10) Axis Design Tool bandwidth estimate.21
Figure 11) Physical and virtual machines sizing example.22

1 Introduction

NetApp® E-Series storage arrays provide performance, efficiency, reliability, and enterprise-class support for large-scale video surveillance deployments.

All video surveillance management software shares the common feature of recording live video feeds to storage for subsequent replay to aid in forensic analysis or investigation of persons or events within the field of view of a single camera or group of cameras. These video feeds, generated by hundreds or thousands of cameras, are typically configured to record continuously, 24 hours per day, 7 days per week, with retention periods in the range of months to years.

1.1 Publication Scope

This document is intended to provide a set of guidelines for sizing a video surveillance solution based on NetApp E-Series storage arrays.

1.2 Audience

This publication is intended to provide guidance to physical-security integrators, video surveillance management software engineers, network and storage system engineers, and architects responsible for integrating NetApp E-Series storage systems into existing video surveillance deployments or designing and implementing new deployments.

The content in this report is presented with the expectation that these professionals can use this information, combined with their experience and supporting documents, to build an efficient, scalable, and highly available system.

Targeted Deployments

The targeted deployments for this introduction are large (200–1,000 cameras or more) with retention periods of at least 30 days and primarily use HDTV/megapixel resolution cameras.

1.3 Training Offerings

There are a number of web-based and instructor-led training opportunities to enable successful deployment of the NetApp E-Series storage array. The classes listed in the [NetApp University Customer Learning Map](#) under Storage Systems are recommended end-user training classes.

Table 1) Training offerings.

Class Description	Duration (Hours)	Delivery
E-Series E5400 Technical Overview	01:00	Web-based
E-Series E2600 Technical Overview	01:00	Web-based
NetApp E-Series Hardware Architecture and Configuration	00:45	Web-based
Configuring NetApp E-Series Storage Systems	24:00	Instructor-led
Maintaining NetApp E-Series Storage Systems	16:00	Instructor-led

2 Sizing Considerations

Video surveillance solutions based on NetApp E-Series provide performance, efficiency, and reliability with enterprise-class support for large-scale video surveillance deployments. The solutions can utilize the

NetApp E-Series storage array in an E5460 configuration or an E2660 configuration. This section addresses sizing guidance for both deployment models. A variety of video surveillance resources are available on the [Field Portal](#).

2.1 System Requirements


The system requirements are specified in a request for proposal/quote developed by either the end customer or a physical-security consultant in contract with the end customer.

The physical-security integrator must work with the physical-security manager to accurately assess specific requirements, including:

- Retention period
- Number, location, and type of cameras, resolution, frame rate, and so on
- Video management software selected
- Number of cameras per recording server
- Continuous recording or record on motion
- Frequency of viewing archived video
- Failover design requirements

These requirements have dependencies that affect the total system and are illustrated in Figure 1.

Figure 1) System requirements.



The individual requirements must be discovered, analyzed, and documented to accurately size the storage array. The following sections address each of the components, provide technical background on the options, and make recommendations on industry best practices.

2.2 General Considerations

The process of sizing storage to address the video management application requirements is an exercise that balances throughput and capacity considerations. In video surveillance deployments that are characterized by minimal forensic analysis of archived video, capacity is the primary consideration. Recording servers configured to continuously record video exhibit a relatively deterministic I/O pattern. The arrival rate of video feeds from IP cameras is constant, and the video recording server in turn writes these streams to an archive at a consistent rate. Many video management systems write the video archive to a temporary or recording directory on the volume logical unit number (LUN) and subsequently read and write the temporary video archive to a permanent directory. Although this movement of video data files from one location to another adds a secondary I/O to the initial write, the workload characteristics can be quantified.

Markets that fall into this category are secondary and higher education, enterprise, and commercial or retail deployments. Video is only retrieved and analyzed if an incident requires investigation. For example, the education market might only review video from a few cameras several times a week. The I/O is mainly writes, and there is little viewing (reads) of the archived video.

The gaming market is an example where there is a high degree of forensic analysis of video archives. In these deployments, several security operators are frequently reviewing video archives to investigate suspicious activity, theft, and fraud. This activity adds a high degree of read I/O to the workload, which is exaggerated by the use of fast forward capabilities of the video management software (VMS). These I/O patterns are less deterministic than the constant influx of video feeds from camera to the server because the workload is a function of the number of operators conducting investigations, the number of camera archives being viewed, and the playback speed.

These deployments might require the use of two volumes (LUNs) per recording server. One volume is often configured for RAID 10 for the recording directory, and a second volume for storing the video archive to the target retention period is configured as RAID 5 or RAID 6. RAID 10 is used for the recording directory due to better assumed or observed read performance than RAID 5 or RAID 6 on the storage array. However, if the performance required by the application is achievable with a recording directory using RAID 5 or RAID 6, then this dual (tiered) volume approach adds to the cost and complexity of the implementation, with no functional advantage.

In many cases, the measurable differences between RAID 10 versus RAID 5 or RAID 6 performance might be trivial in a well-designed solution, and the advantage of RAID 10 might be anecdotal from an older, atypical deployment using a poorly performing storage array.

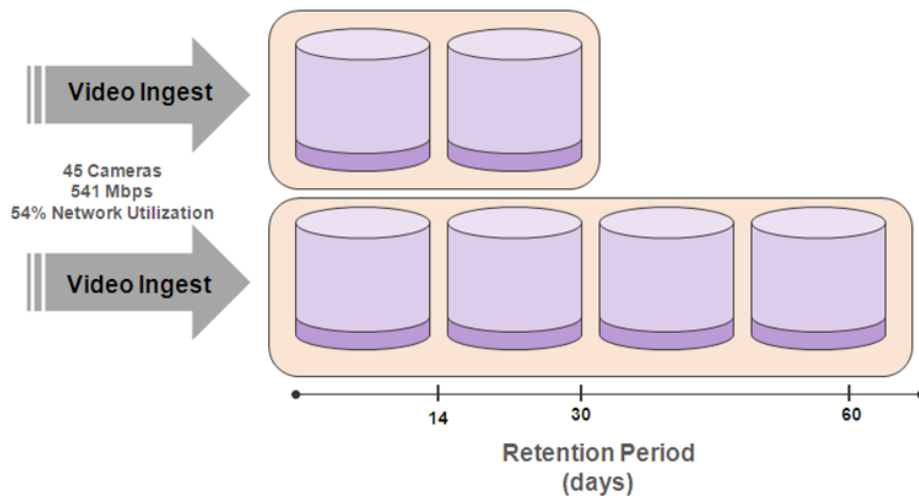
2.3 Retention Period

The retention period is usually determined by organizational or regulatory policy. For example, the Nevada Gaming Commission Regulation 5.160(2), "Surveillance Standards for Nonrestricted Licensees," specifies a minimum retention period of 7 days. However, cameras deployed in nonregulated areas such as lobbies, parking lots, and other common areas might have a 30-day retention period governed by hotel policy.

State and local governments typically specify video retention periods as part of their state records retention schedules. Geographies with a history of criminal activity or terrorist threats might specify longer retention period requirements.

Increasing the retention period does not increase the arrival data rate of video feeds from IP cameras to the recording server and storage array. Longer retention policies make the sizing exercise more of a capacity calculation than a performance consideration. This is illustrated by the example shown in Figure 2.

Figure 2) Example of throughput versus retention.



There is an absolute and a practical limit to the amount of storage allocated to an individual recording server. Additionally, factors such as network interface capacity and CPU, memory, and application performance also limit the number of cameras per recording server. In most video surveillance deployments, the performance characteristics of the E-Series controllers meet or exceed the requirements of the video management application. All video management systems reach a steady state, where video archives are deleted at the same rate that new video files are added to the storage array.

NetApp recommends a retention period that meets or exceeds the applicable policy.

2.4 Reserve Capacity

When the retention period policy is defined, the amount of reserve capacity must be considered. As a best practice, each volume should be maintained at approximately 80% utilization. Most VMS packages implement a file deletion trigger when the configured maximum size of the archive is reached. Each volume owned by the recording server has a configurable threshold defined as the minimum free space or maximum archive size. Usually this is specified in gigabytes/terabytes in addition to the configured (days) retention period. If either of the two maximum values is reached, that is, the retention period in days or the maximum archive size, the oldest video files are removed from the system to prevent a disk full error condition.

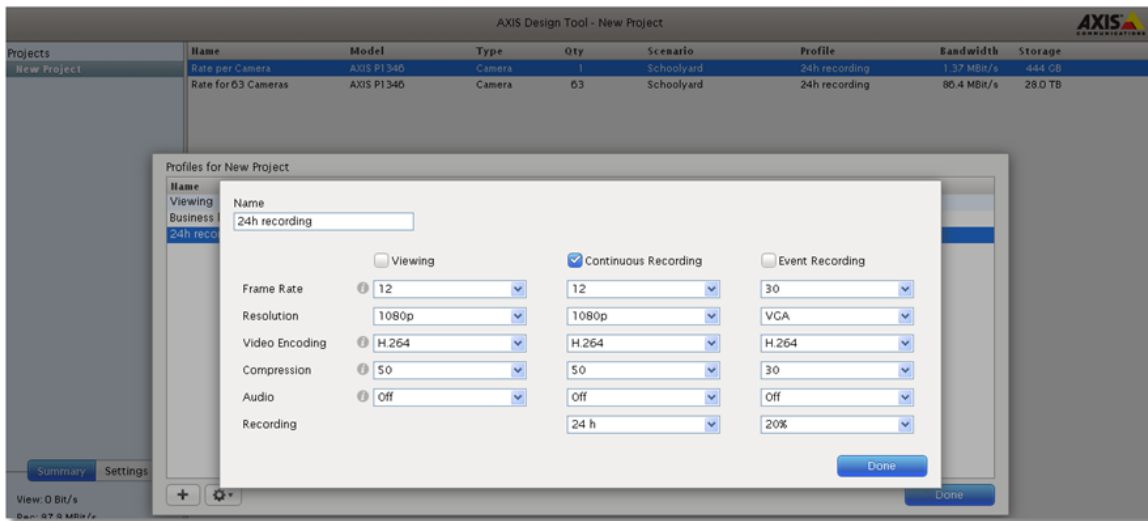
As a best practice, size the system based on a target capacity of 80% based on the stated retention period of the organization. Then, when implementing the system, configure the VMS to use all the configured capacity, maintaining only 5% unused space. This method makes sure that the video retention meets or exceeds the specified policy.

For example, if the target retention period is 30 days, during implementation specify a 38-day retention period with minimum free space at 5% of capacity. This method forces the VMS to delete files based on capacity, utilizing 95% of the volume rather than 80% of the volume.

2.5 Cameras

The number of cameras and the configured resolution, frame rate, codec type, compression factor, and image complexity must be determined to estimate the aggregate video data rate and the amount of storage required. Most camera manufacturers provide a design tool to estimate the data rate and storage requirements based on the camera model and specified values. Figure 3 illustrates an example of one such tool.

Figure 3) Axis Design Tool.



For more information on the Axis Design Tool, refer to the [Axis](#) website.

These tools provide only an estimate. The results might vary and should be verified through field trials run by the physical-security integrator.

The total number of cameras at a specific location is a function of the physical-security requirements of the site and how the security integrator plans to address these requirements. The total population of cameras will encompass a variety of camera models, and in some cases the cameras might be from different manufacturers. The integrator can select from several types of cameras, including indoor or outdoor cameras; fixed or an, tilt, zoom (PTZ); and tamper-proof or vandal-proof dome cameras. From a sizing perspective, the type of camera is not important, but the resolution and number of channels (video feeds) from each camera are important.

When the total number of cameras is determined, they need to be further categorized by the resolution.

Resolution

There are a wide variety of resolutions available in the video surveillance camera industry. They are distinguished by these categories:

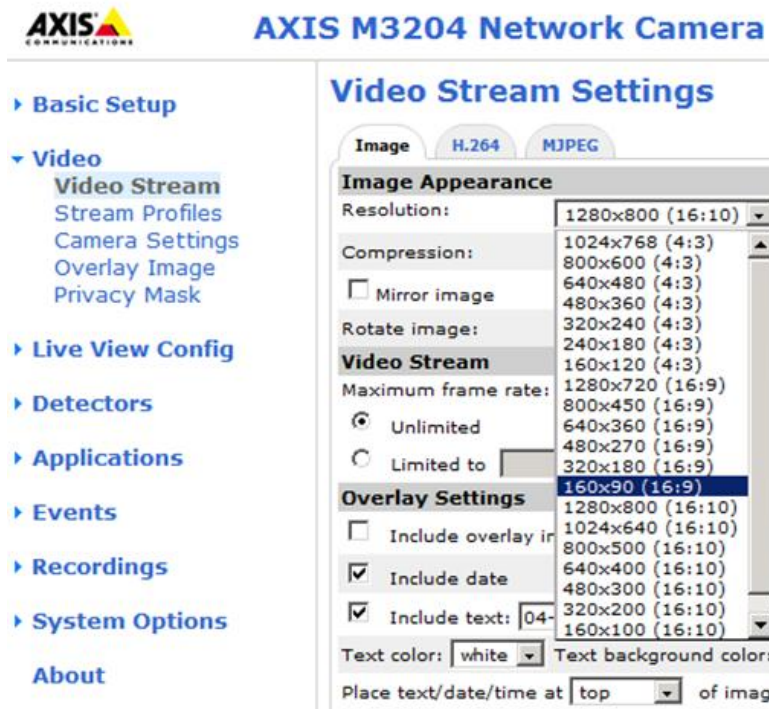
- Analog video (NTSC/PAL); 4CIF is commonly used: 704x480 pixels or 0.4 megapixels
- Video graphics array (VGA) to XVGA: 640x480 pixels to 1024x768 pixels up to 0.75 megapixels
- Megapixel SXGA to QSXGA: 1280x1024 pixels to 2560x2048 pixels or 1.3 to 5.2 megapixels
- High-definition television (HDTV): 1280x720 pixels and 1920x1080 pixels or 0.9 to 2 megapixels

E-Series storage is ideal for HDTV and megapixel resolution deployments because of the high density and performance characteristics of the E-Series. For new installations, HDTV/megapixel cameras are the preferred choice over standard definition (SD) cameras. Although the purchase price of a HDTV/megapixel camera is slightly more than that of an SD camera, the installation cost is the same. The total cost of cameras and installation might be less for HDTV/megapixel deployments over SD cameras, because fewer cameras are required to effectively cover an area.

Video surveillance camera models are selected to meet a functional requirement. These requirements are classified as detection, recognition, and identification. There are industry guidelines for the number of pixels per foot required to address each category. Detection has a lower resolution requirement than identification. One HDTV/megapixel camera might provide sufficient resolution to meet the pixel-per-foot requirements, where earlier multiple SD cameras would need to be deployed.

Network video cameras can be configured at resolutions below the specified maximum resolution. An example of the configurable resolutions for the Axis M3204 network camera is shown in Figure 4.

Figure 4) Video stream settings.



An HDTV/megapixel camera can also be an SD camera, an HDTV format camera, or a megapixel format camera. The Axis M3204 camera can be configured for HDTV format (1280x720 resolution, 16:9 aspect ratio), megapixel resolution WXGA (1280x800 resolution, 16:10 aspect ratio), or VGA (640x480, 4:3 aspect ratio) resolution.

NetApp recommends HDTV/megapixel cameras for new deployments.

Frame Rate

The configured frame rate of the network video cameras also must be determined to accurately estimate the required storage capacity. Frame rates in the gaming industry are specified by regulations and are required to be 30 frames per second. Common frame rates for other industries are 7 to 12 frames per second or less. Cameras positioned at cash registers and teller stations usually require at least 12 to 15 frames per second. In school or office hallways, 5 frames per second is usually sufficient. Parking lots and other overview scenes for detecting cars, people, or objects often require only 1 to 3 frames per second.

Cameras positioned with horizontal movement across the field of view or with high-speed movement (highway intersections, for example) generally require higher frame rates than scenes with vertical movement or slow-moving people or objects. As a reference, motion pictures (35mm sound film) have traditionally used 24 frames per second. The human eye begins to notice choppy motion below 16 to 18 frames per second.

The specified frame rate greatly influences the network bandwidth and storage requirements. In addition to categorizing the cameras by resolution, now frame rate is also added to the equation.

NetApp recommends using a frame rate that meets the regulatory or functional requirements of the camera.

Compression Type

The compression type (compression standard) options for video surveillance deployments are Motion JPEG or MPEG-4/H.264. Motion JPEG is a series of individual images that have no interdependency between frames. Motion JPEG is often required by analytic software implementations. Because there are no interframe dependencies, Motion JPEG is used in networks that exhibit rates of packet loss that would make MPEG-4/H.264 unusable. As the frame per second rate increases, the bandwidth and storage requirements for Motion JPEG become increasingly costlier compared to MPEG-4/H.264 for a given resolution.

MPEG-4 (MPEG-4 Part 2) and H.264 (MPEG-4 Part 10 AVC) are compression standards that transmit a reference frame periodically and send the changes in the scene in subsequent frames. The frequency of reference frames is configurable by the value specified for group of video/group of pictures (GOV/GOP) length. H.264 is more efficient than MPEG-4, and its use has superseded MPEG-4 for HDTV/megapixel video cameras.

Typically, motion JPEG uses Transmission Control Protocol (TCP) to transport video between a camera and a server. TCP is connection-oriented and provides for reliable transport. Alternately MPEG-4 and H.264 are usually transported in Real-time Transport Protocol (RTP)/User Datagram Protocol (UDP). The UDP transport is connectionless and does not retransmit lost packets. The TCP protocol favors reliability over timeliness. The majority of voice and video applications use RTP/UDP transport.

The network infrastructure must exhibit a very low packet loss for transport of MPEG-4/H.264 for acceptable image quality. Dropping as little as 1/20th of one percent of the IP packets for a H.264/RTP/UDP video stream is noticed as trailing artifacts in the image. The distortion of the image is corrected when the next reference frame is successfully received.

In video surveillance deployments, NetApp recommends using H.264 and RTP/UDP transport for network bandwidth and storage efficiency as a best practice.

Compression Ratio

The compression ratio is a configurable value on the network video camera that specifies to what factor the image is reduced, before transmission to the recording server. A 10% value indicates little compression, and a 90% value indicates a high degree of compression. The process of compressing the image is known as quantization. In video image processing, this is implemented as lossy compression, meaning details are lost to reduce the size of the image. Selecting a high value might result in an image with compression artifacts or pixelization. An appropriate value depends on the scene complexity, lighting, shapes of objects, and colors in the scene. Values of 30% to 50% are common.

NetApp recommends leveraging the recommendations of the camera manufacturer and experience of the system integrator when selecting an appropriate compression ratio.


Variable or Constant Bit Rate

The H.264/MPEG-4 video encoder in the network video camera may be configured for either variable or constant bit rate. Constant bit rate (CBR) varies the image quality to maintain a constant output network bit rate. If there is little motion in the scene, the quality will remain high. If there is a complex scene with motion (for example, trees swaying in the breeze), the image quality decreases. The decrease in image quality is recognized as noticeable pixilation of all or part of the image.

Variable bit rate (VBR) maintains the image quality but changes the output network bit rate to accommodate motion in the scene. Network video cameras deployed in areas with little or incidental motion will have lower bandwidth and storage requirements with VBR. In low-light settings, the camera imager might introduce image noise, which has the same effect as motion. The sizing calculator tool of the camera manufacturer might allow the user to specify level of detail and percent motion or select an image scenario such as intersection, stairway, and so on. A night option that will return higher bandwidth and storage requirements compared to the daylight version scenario might also be available. Figure 5

illustrates the effect to ambient indoor lighting levels being adjusted at 7 a.m. through 7 p.m. with 64 HDTV cameras recording onto a single volume (LUN). In this example, there is approximately a 7% difference between the two data rates and storage required.

Figure 5) Daylight versus nighttime data rate.



It is important to accurately estimate the selected scene complexity because it greatly influences the bandwidth and storage required. In some cases, the storage required for a scene at night might be 60% higher than during the day. Physical-security integrators use infrared (IR) illuminators to provide additional light in low-light environments. The use of IR illuminators minimizes video image noise and the resulting increase in bandwidth and storage.

NetApp recommends using VBR for H.264/MPEG-4 and accurately estimating the scene complexity and percent of motion.

Audio

Most network video cameras also support an audio channel. The network bandwidth and storage required for audio are minimal compared to video. When sizing, estimate approximately 30Kbps if the audio channel is stored with the video.

Video Management System

The VMS software is the platform for managing the network video cameras, processing of video streams from the cameras, and recording of streams to the storage array for archive and retrieval. The VMS software also integrates with access control systems, manages events, and generates alerts. Failover recording servers can be configured to continue archiving video streams in the event of a hardware failure.

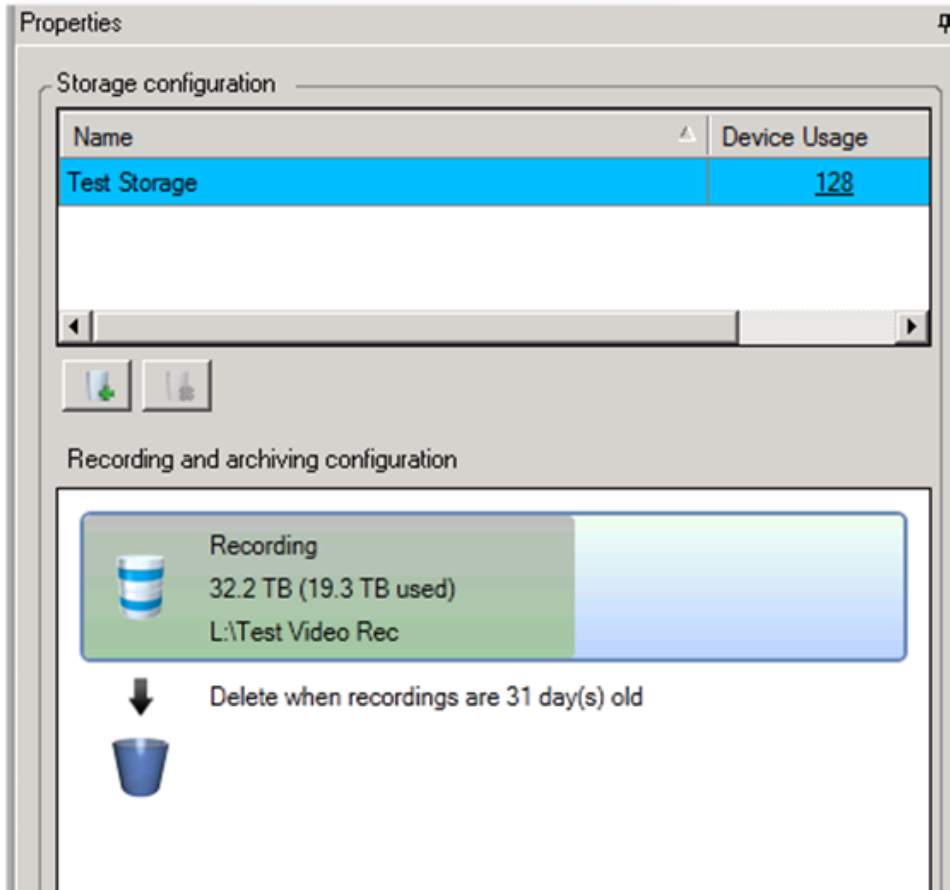
Feature Set Limitations

Video management systems might be licensed by a feature set. For example, OnSSI Ocularis has four feature sets: PS, IS, CS, and ES. The ES feature set supports thousands of cameras at multiple locations and utilizes a single central management platform for all recording servers. There are no fixed limits to the number of cameras per recording server, but the documented recommendation for HDTV/megapixel cameras is 40 per recording server instance. There is no absolute limit on the number of days of retention

on the live recording volume (LUN). The distinction between multiple recording locations is further explained in the section on Tiered Storage.

Following is an example of a single recording location. Figure 6 illustrates Ocularis ES configured for 31-day retention in the live recording volume (LUN), for 128 cameras on a single recording server.

Figure 6) OnSSI Ocularis storage configuration.



Alternately, the Ocularis CS feature set has an absolute limit of 64 cameras per recording server, and the live recording volume is limited to seven days of recordings. Each recording server is configured independently; there is no central management feature.

The selection of the appropriate VMS and feature set is a collaborative effort between the video security integrator and the end user. Regardless of the VMS software package selected, any performance or hard limitations of the number of cameras per server or absolute size of volumes must be assessed to accurately size and configure the storage array.

Number of Volumes (LUNs) per Server

Most video management software can write video archives to multiple drives. For example, the recording server might have both an E:\ and F:\ drive as a target location for storing video archives. The recording server might have limitations that prevent a single camera from being written to separate drive letters. The recording server will typically select the location with the most available space when making a determination between drives.

It is more efficient to use a single volume rather than several volumes because a minimum amount of free space is required on each volume.

Continuous Recording or Record on Motion

VMS can be configured to record continuously or to record a length of time before and after a triggered event. Record on motion is the most common event to trigger recording. Because video cameras are commonly placed in areas that routinely have little motion (for example, emergency exits) or have little motion for some period of time (for example, school hallways at night), implementing record on motion can greatly reduce storage requirements.

Although configuring record on motion is efficient in the use of storage, it also has several disadvantages. Implementing server-side record on motion is CPU intensive and reduces the number of cameras that a recording server can support compared to continuous recording. Record on motion does not reduce the throughput requirements of the storage array because the recording server must temporarily store the video to disk to implement prebuffering. Prebuffering is the amount of video retained before the event. The amount of video to retain following the event is also configurable. Many video management systems store video in files that contain three to five minutes of video from each camera. It is common to implement record on motion by deleting all files except the files that contain the prebuffer and postbuffer of the event.

Record on motion also can be complicated and time-consuming to configure. Usually a detection area in the scene is selected to be the target of the record on motion algorithm, for example, a doorway. Then the degree of sensitivity is selected. A change in lighting also triggers a motion event. The algorithm detects changes in either the chroma (color) or luminance (brightness) of pixels. If the sensitivity is set too low, subtle changes do not trigger the motion event, and the video is not archived. If the sensitivity is set too high, the disk usage approaches that of continuous recording.

Other factors such as object size and percentage of object in view might also require configuration.

VMS might support camera-side motion detection with little CPU overhead on the recording server. In a camera-side implementation, the CPU on the network video camera executes the motion detection algorithm. When motion is identified, an alert is sent to the recording server to indicate that motion was detected. The server retains the video files that contain motion and deletes files that do not. Camera-side analytic algorithms might also support video analysis such as tripwire, people counting, objects left behind, or individuals loitering events.

The NetApp recommendation is to evaluate the tradeoff between reducing storage costs against the increase in system management costs when implementing record on motion.

Video Walls


Video walls are one or more client workstations and monitors used for displaying video in a control room setting. Video can be pushed to the workstation running in a video wall mode and might also display video as the result of triggered events. From the sizing perspective, video displayed on a video wall requires the same amount of throughput as a client-viewing workstation.

Viewing Archived Video

The frequency and number of cameras archives viewed concurrently will incur read I/O to the storage array. Video archives may be viewed at the normal speed or at an increased playback speed. The application might support increasing the playback speed over one thousand times the recorded speed.

Figure 7 demonstrates the performance characteristics of viewing 64 HDTV cameras at normal playback speed and transitioning to 16x playback speed.

Figure 7) 64-camera transition from 1x to 16x.



In this example, the average data rate increased from approximately 140Mbps (2.2Mbps per camera) to 400Mbps (6.2Mbps per camera) when transitioning from 1x to 16x playback. Note that the data rate increase is not a linear relationship; a 16x speed increase increased the I/O rate by a factor of 3x.

The performance characteristics associated with the forensic capabilities of the VMS-viewing client are implementation-specific and might vary between releases. Both video walls and client-viewing workstations affect the performance characteristic of the storage array but not the capacity required of the storage array. Investigative activities by the client workstations and incident reporting files might need to be considered when sizing the storage array.

2.6 Centrally Stored Video Clips

Centrally stored exported video clips (OnSSI Ocularis refers to these as bookmarks) may be stored on volumes (LUNs) on the storage array. The location is configured during the software installation.

The amount of storage required is dependent on the frequency and size of the video archives stored. These files have an infinite retention period. They are only deleted manually by the operator.

2.7 Tiered Storage

Tiered storage is the term for distinguishing between two or more types of storage that vary in price, performance, capacity, or function. Several VMS packages implement an archival function that moves video files from a temporary or recording location to an archive location. This is referred to as tiered storage because the function aspect of the data has changed.

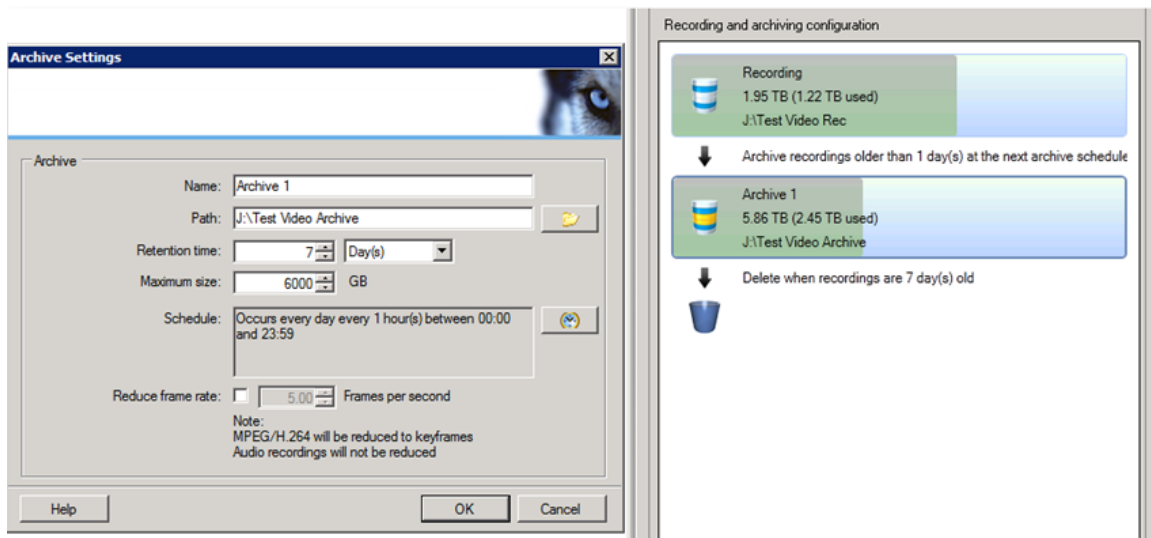
Note: Industry documentation often refers to storing video files in a database. This terminology does not imply that the video is stored in a relational database structure; rather, the individual files are stored within a directory structure of the file system.

Archiving is the term used to describe the function of moving video files from their original location on disk to an alternate location. This configuration is defined on a server-by-server basis. The configuration might have zero, one, or more archive locations. These archive locations can be defined on the same volume (LUN) as the recording location or on a separate volume (LUN) on the same or a different storage array. For example, the recording location volume (LUN) might be RAID 10, whereas the archive location could be RAID 6.

When using multiple storage arrays, they need not be the same type; block-level SAN storage can be the source of the archive, and a file-level NAS device can be the destination. Files in the archive location may be backed up to tape or other media provided the archive process is not active.

During the archive process, video files may be groomed to reduce the frame rate of the original recording or may be encrypted. The archive process occurs on a user-configured schedule. It might be hourly, every four hours, or daily. This archive process affects both the performance and capacity of the storage array. Figure 8 is a sample configuration of separate directories on the same volume (LUN) for recording and archiving.

Figure 8) OnSSI Ocularis ES recording and archiving configuration.



In the example, both maximum size and retention time are configurable. Files in the recording directory are moved to the archive directory when the files are over 24 hours old. The archive process runs every hour. They are stored in the archive directory for seven days and then deleted.

Drive letter J: is mapped to a RAID 5 volume group with 10.913TB of capacity. The recording server will not use all of the available capacity on drive J: because of the maximum size parameters specified on the recording (1.95TB) and archive (5.86TB) configuration. This configuration can use up to 7.81TB of the 10.913TB available, or approximately 72%.

Performance of Tiered Storage

During the normal processing of video feeds from cameras and writing to the video recording storage location, the data rate is relatively constant between the server and storage array. When the archive process is initiated, the I/O characteristics change, based on the source and destination of the archive. If the administrator has configured separate volumes (seen by recording servers as LUNs) for the recording location and archive location (for example, E:\RECORDING and F:\ARCHIVE), then the volume containing the recording location will incur read I/O, and the archive location will incur write I/O. The archive function in this example changes the I/O characteristics of the recording location from primarily writes to both reads and writes for the duration of the archive process.

The duration of the archive process is a function of the amount of data that must be moved. It is important to understand that the recording server writes video to storage at approximately the rate of arrival from the networked video cameras. The archive process will read and write data as rapidly as the recording server can read, process, and write the files to the destination. In testing, the I/O rates have been observed to increase by eight times or more during the archive process.

Sizing of Tiered Storage

Every recording location and archive location must have some free space for the system to record arriving video streams. If the locations have insufficient space, the oldest video will be deleted or autoarchived if possible. If files cannot be autoarchived or deleted quickly enough to reclaim space, the recording server might not be able to write video streams to disk. It is important to allocate sufficient capacity to the respective archive locations to meet the requirements of the video retention policy. Additionally, the archive process must occur more frequently than the configured retention period for the location. For example, if the retention period is seven days, the archive function should occur at least once a day; every four hours or hourly is also commonly implemented.

2.8 High Availability

A failover server is an idle recording server that can assume the recording server role in the event a primary recording server becomes unavailable. This enhances the availability of the system; however, video is lost during the time taken to detect the server failure and initiate recording on the failure server.

The implementation deploys a keep-alive (heartbeat) TCP session between the primary and failover servers. Failover servers are defined in a group, and the primary server is configured to specify the failover group. This allows for greater flexibility and efficiency because there is no one-to-one relationship between the active server and failover server. For example, a design utilizing 10 recording servers might be configured with 2 or 4 failover servers.

Sizing of Failover Volumes

The failover servers must have sufficient disk capacity to assume the archiving function for any of the recording servers referencing the failover group. The volumes (LUNs) defined to the failover server are not used unless the failover event is triggered. When the failed server is restored and back online, the video archives on the failover server are migrated to the primary server. This recovery process adds additional I/O to the primary server and its LUNs during the migration process.

The volumes (LUNs) defined to the failover servers must be of a sufficient capacity to store video for the length of time required to detect that a primary server has failed, resolve or replace the failed component, and bring the primary server online. This length of time varies based on the geographic location of the servers, availability of system administrators at the site, and network management practices implemented.

NetApp recommends sizing the volumes (LUNs) of the failover servers to retain a minimum of three to five days of video from the recording server with the largest volume (LUN). To calculate the size of the failover volumes, divide the capacity of the primary volume by the site retention period and then multiply by the number of days to retain during failover. For example, given a 30-day retention, a volume (LUN) of 22TB ($22/30 \times 5$) requires a 3.6TB volume (LUN) for the failover server.

Because of the relatively small size of the failover volume (LUN) when compared to the size of the primary recording volume (LUN), creating a volume group for each failover volume (LUN) is inefficient. Allocating one volume (LUN) per volume group eliminates disk-head contention when constantly writing to separate volumes (LUNs) in the volume group. In the example where a 3.6TB volume (LUN) was required for the failover server, creating a separate volume group for each failover server would require the following number of drives to meet the capacity required:

- 4 drives: RAID 10 (5.5TB)

- 3 drives: RAID 5 (5.5TB)
- 5 drives: RAID 6 (8.3TB)

Failover volumes (LUNs) are idle unless a failover server is engaged. They are used only if a recording server is taken out of service for an upgrade or if a failure occurs. NetApp recommends creating a volume group or disk pool for failover volumes and allocating volumes (LUNs) to meet the requirements of the failover servers.

As an example, a disk pool with 12 drives provides for 25.644TB of capacity. From that capacity, five volumes can be created, each with a capacity of approximately 5TB.

Volume	Capacity	DA	Enabled
POOL_2_FAILOVER_VOL2A	5,251.071 GB	No	No
POOL_2_FAILOVER_VOL2B	5,251.071 GB	No	No
POOL_2_FAILOVER_VOL2C	5,251.071 GB	No	No
POOL_2_FAILOVER_VOL2D	5,251.071 GB	No	No
POOL_2_FAILOVER_VOL2E	5,252.000 GB	No	No

Using a dynamic disk pool (DDP) for failover volumes provides a RAID 6 level of protection and inherent hot spare coverage with an efficient use of disk drives.

Multicast and Secondary Streams

An alternate method to achieve high availability is to use multicast and secondary streams. Most IP cameras support the transport of video streams as IP multicast packets. IP multicast is a bandwidth conservation technique where the network devices replicate these packets and deliver them to multiple receivers. Some video management systems can be configured to define a single IP camera for IP multicast transport and receive the video streams on two or more recording servers. Each recording server archives the video independently. This is a means to maintain two or more copies of a video feed from a single IP camera. It is an alternative means of providing high availability without the need for primary and failover servers. Not all video management systems support this configuration.

Another means of recording the same images from an IP camera is through the use of secondary video streams. Most IP cameras have the ability to unicast a primary video stream to one recording server and have a second recording server access a secondary stream. In some instances the secondary stream is at a lower frame rate, resolution, or compression algorithm than the primary stream. For example, H.264 might be supported as the primary stream, but M-JPEG is only supported on the secondary stream.

When deploying cameras overlooking scenes of critical importance, using two or more cameras with overlapping field of views and defining these cameras to separate recording servers is a best practice for high availability.

The implementation of multiple unicast video streams or IP multicast transport is an alternative to providing high availability without the need for primary and failover video servers.

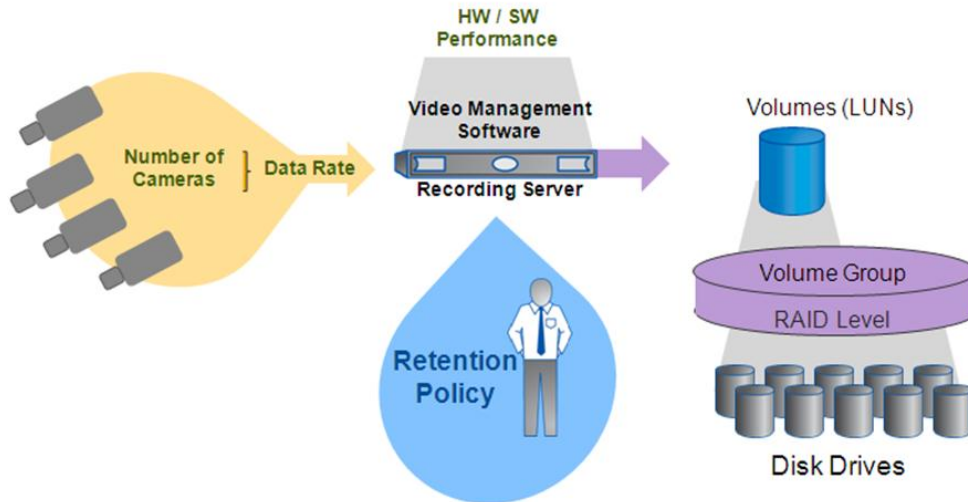
3 Sizing Fundamentals

To properly size the storage array, three fundamental questions must be answered:

- What is the data rate of video received per server?
- What are the number of servers required?
- What is the retention policy?

After the data rate per server is determined, the size and number of volumes (LUNs) required by each server can be calculated based on the retention period. This process is illustrated in Figure 9.

Figure 9) Sizing fundamentals.



If all the cameras in the deployment are configured identically, the data rate is the same for all cameras, and the discussion can be simplified as a number of cameras per server. Each server has the same number of cameras, with the resulting data rate at or below the specified maximum for the server. An example of this sizing methodology is described in the Sizing Considerations section.

The VMS vendor documentation should provide a recommendation for performance and scalability. The number of cameras supported per server is a function of the version of the VMS system, the performance characteristics of the server, and the average data rate of the cameras.

Many video surveillance deployments, however, encompass a variety of camera models at different resolutions, frame rates, compression ratios, and image complexities. In such deployments, the sizing exercise becomes more complex due to the increased number of variables.

3.1 Storage, Operating System, and File System Capacity Considerations

E-Series

Video surveillance solutions based on E-Series are configurable with either an E5460 or E2660 controller configuration. Either traditional volume groups or DDP is supported. The supported disks are 3TB NL-SAS (4TB drives will soon be available) and 900GB 10K SAS.

Performance for both traditional volumes and DDP on a volume-by-volume basis is more consistent if there is one volume per volume group or disk pool. Provisioning the storage array in this manner eliminates disk drive head contention when compared to provisioning multiple volumes (LUNs) per volume group.

If using traditional volume groups, NetApp recommends allocating one hot spare drive for every 30 disk drives. A 60-drive shelf would therefore need at least two hot spare drives. DDP does not require a separate hot spare drive to be allocated. Spare capacity is reserved on all the drives in the disk pool to provide for high availability in the event that a disk drive in the pool fails. Table 2 provides a reference for sizing the solution.

Table 2) Usable capacity by RAID level.

Number of Disks	Dynamic Disk Pool (RAID 6): 3TB	RAID 6: 3TB	RAID 5: 3TB	RAID 10: 3TB	RAID 10: 900GB
2				2,794.019GB	837.863GB
3			5588.039GB		
4			8382.059 GB	5,588.039GB	1675.7GB
5		8382.059GB	10.914TB		
6		10.914TB	13.642TB	8,382.059GB	2513.6GB
7		13.642TB	16.371TB		
8		16.371TB	19.099TB	10.914TB	3351.5GB
9		19.099TB	21.828TB		
10		21.828TB	24.556TB	13.642TB	4189.3GB
11	21.375TB	24.556TB	27.285TB		
12	21.371TB	27.285TB	30.013TB	16.371TB	5027.2GB
13	23.511TB	30.013TB	32.742TB		
14	25.648TB	32.742TB	35.470TB	19.099TB	5865GB
15	27.789TB	35.470TB	38.199TB		
16	29.925TB	38.199TB	40.928TB	21.828TB	6702.9GB
17	32.062TB	40.928TB	43.656TB		
18	34.199TB	43.656TB	46.385TB	24.556TB	7540.8GB

RAID 10 or disk mirroring and striping is configured by selecting RAID 1 with four or more drives. The usable capacity of a 3TB drive is 2794.019GB. The gray cells in the preceding table are unsupported selections.

The maximum number of disk drives in a volume group is 30. The minimum number of drives in a DDP is 11; the maximum number of disks in the pool is the total number of drives in the array.

After calculating the number of disks required supporting the volumes (LUNs) required for all servers in the deployment from Table 2, use Table 3 to determine how many controller shelves and expansion shelves are required.

Table 3) E-Series disk shelves for video surveillance deployments.

Category	E5460	E2660
Form factor	4U/60 drives	4U/60 drives
Maximum disk drives	360	180
Controller shelf	1	1
Maximum expansion shelves	5	2

Category	E5460	E2660
Total (maximum) number of disk shelves	6	3

The following example illustrates this process. A deployment for 640 cameras, 64 cameras per server, and 10 servers is assumed. Each camera on an average generates a 1.8Mbps video feed. The physical-security integrator has determined that the number of 3TB disks required to meet the recording, archive, and failover databases for 30-day retention with one hot spare for each of the 30 drives is 231. Because the E2660 can only support 180 drives, the E5460 with a DE6600 controller shelf and three expansion shelves can house 240 disks.

For this deployment, the E5460 would be the recommended choice. The alternative would be to deploy two E2660 chassis, each with one expansion shelf.

New Technology File System

New technology file system (NTFS) is the preferred file system for Microsoft® Windows® operating systems. The majority of VMS systems are built on Windows Server® 2008 R2 or later.

Many VMS software packages recommend formatting disks with an allocation unit size of 64KB. Given this recommendation, the maximum NTFS volume size is approximately 256TB.

Most VMS systems store video from a single camera to a file with either a maximum size or a maximum number of minutes of video. These parameters might or might not be configurable by the end user. Typically, the size of the file is 3 to 5 minutes of video from a single camera or in some cases up to 30 minutes. The documented maximum NTFS file size is approximately 16TB, which poses no practical limitation.

VMware ESXi 5.1

Implementations using VMware® ESXi™ 5.1 have a 2TB volume size limit for VMFS-5 datastores.

Based on the number of cameras per server and retention period of most deployments, a recording server running in a virtual machine will typically need a volume (LUN) in the 10TB to 30TB range. VMware raw device mapping (RDM) provides access to a volume (LUN) on the storage array. RDMs require the mapped device to be a whole LUN. The documented maximum RDM volume (LUN) size is 64TB.

4 Video Surveillance Sizing Examples

This section illustrates several sizing scenarios. They increase in complexity to show various situations and how they may be sized by the video surveillance integrator.

4.1 Sizing Example 1: A Simple Deployment

The first example is for a base configuration for a video surveillance solution provisioned with two physical servers hosting four virtual machines and one E2660 shelf with 60 3TB disks. This sizing example assumes there are no failover servers implemented and the video management software runs under Windows 2008 R2.

The sample deployment assumes 64 Axis P1346 cameras per server for a total of 256 cameras in the deployment. The Axis Design Tool is used to estimate the bandwidth and storage required for this model of camera. In Figure 10, a per-server estimate is shown, with line items for a single camera and then the additional 63 cameras to allow the report to calculate the totals.

Figure 10) Axis Design Tool bandwidth estimate.



AXIS Design Tool
BANDWIDTH ESTIMATE

Project: New Project Date: Tuesday, August 21, 2012 Storage: 30 days

Name	Model	Qty	View	Rec	Event	Storage
Rate per Camera	AXIS P1346	1	0 Bit/s	1.37 MBit/s	0 Bit/s	444 GB
Rate for 63 Cameras	AXIS P1346	63	0 Bit/s	86.4 MBit/s	0 Bit/s	28.0 TB
Project Summary		64	0 Bit/s	87.8 MBit/s	0 Bit/s	28.4 TB

Important notice: The AXIS Design Tool is provided for guidance only. The estimates and recommendations produced by the use of AXIS Design Tool results from a limited number of test scenarios and Axis' conclusions thereof. The user of AXIS Design Tool (the "User") acknowledges and agrees that every installation of a camera system is unique and that the actual results measured from the User's system will differ from the estimates produced by AXIS Design Tool. The estimates produced by AXIS Design Tool are thus supplied upon the condition that the User will make its own determination as to the suitability for its purposes. Axis and/or its representatives will in no event be responsible for damages of any nature whatsoever resulting from the use of or reliance upon the results produced by AXIS Design Tool. **NOTE:** The AXIS Camera Station hardware and software recommendations are only valid up to a total of one hundred (100) channels for one system, regardless of the fact that AXIS Design Tool allows more than 100 channels to be specified for a system.

This profile uses a 30-day retention, recording 24 hours per day at 12 frames per second, and 1080p resolution with H.264 at 50% compression. The image complexity scenario is a schoolyard. Note that this deployment requires approximately 0.5TB per camera per month.

Given this estimate, each server requires a volume (LUN) capable of storing 28.4TB. Referring to Table 2, configuring a 14-disk RAID 6 volume (LUN) provides 32.7TB of usable space. This provides a margin of 14% free space (28.4/32.7).

When SANtricity® ES is used to configure one 14-disk RAID 6 volume per volume group, the volume summary and array summary for the storage array would be as follows:

```

PROFILE FOR STORAGE ARRAY: stle2660-33_34 (Tue May 21 11:00:31 EDT 2013)

Number of standard volumes: 4

NAME      STATUS  CAPACITY  RAID LEVEL  VOLUME  GROUP  LUN  ACCESSIBLE BY
VOL_RACK_1 Optimal 32.741 TB 6      VG_1     1     Default Group
VOL_RACK_2 Optimal 32.741 TB 6      VG_2     2     Default Group
VOL_RACK_3 Optimal 32.741 TB 6      VG_3     3     Default Group
VOL_RACK_4 Optimal 32.741 TB 6      VG_4     4     Default Group

SUMMARY-----
Number of drives:      60
Mixed drive types:    Enabled
Current media type(s): Hard Disk Drive (60)
Current interface type(s): Serial Attached SCSI (SAS) (60)

Total hot spare drives: 4
Standby:              4
In use:               0
    
```

There are 56 disks in use with four hot spares. This example does not include volumes (LUNs) for failover recording servers or centrally stored video clips (bookmarks).

In this sizing example, the assumption is made that a single camera model is deployed with the same image complexity, frame rate, and compression factor. Because there is consistency in the video ingress to the server, it is very easy to determine the number of cameras per server virtual machine. There are no failover servers configured in this example.

Many deployments will use a variety of camera models at a variety of resolutions and frame rates. Failover sizing must also be included. The following sizing example illustrates a more complex deployment.

4.2 Sizing Example 2: Larger System with Failover and RAID 10

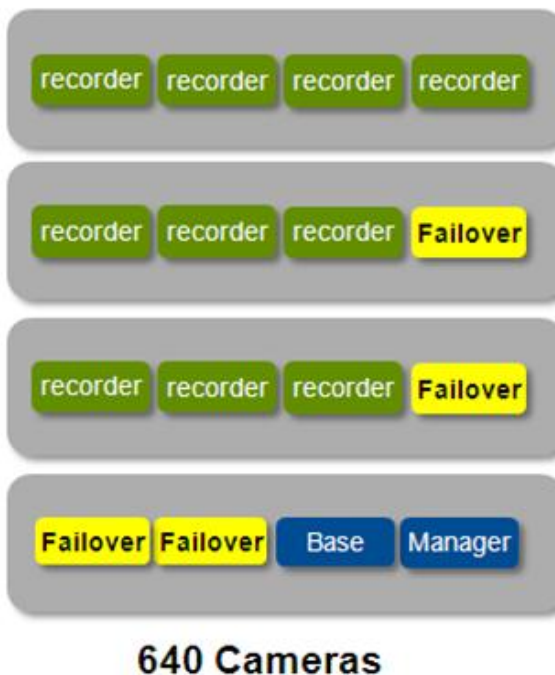
This section illustrates a sizing example that implements sufficient failover recording servers to recover recording of the cameras in the event one physical machine fails or is taken out of service.

This example uses an E2660 controller and two expansion shelves with a total disk population of 150 3TB disks and 30 900GB disks. The VMS is OnSSI Ocularis ES that runs under Windows 2008 R2.

The deployment has a requirement to write the recording database to a RAID 10 volume group composed of 900GB 10K RPM SAS drives. The recording database contains the first 24 hours of recorded video, and the archive database is defined on RAID 6 volume groups composed of 3TB 7200 RPM NL-SAS drives.

The deployment uses the E2660 fully populated configuration with 4 physical servers hosting 10 virtual machines for day-to-day recording and 4 failover virtual machines. The number of cameras per virtual machine is targeted at 64, for a total camera count of 640. This configuration is illustrated in Figure 11.

Figure 11) Physical and virtual machines sizing example.



The proposed camera configuration is Axis models M3204 and P1346 at an HDTV resolution of 1280x720 with H.264 encoding and RTP/UDP transport, 12 frames per second with 30% compression.

As a point of reference, the Axis Design Tool estimates the intersection (night option) image complexity for an M3204 in the planned configuration to be 1.3Mbps, requiring 448.2GB for 31 days of retention. Changing the image complexity to reception area (night option) results in an estimated 726Kbps (231.8GB for 31 days retention).

Given these estimates, the physical-security integrator plans to implement 12+2 volume groups (14 disks) for the archive database and a 1TB volume for the recording database. To accommodate the volumes for

the failover servers, the first four volume groups will contain two volumes: a failover volume and an archive volume. Due to space limitations, the failover servers will not be configured with a separate recording database on the RAID 10 10K RPM drives.

The failover volumes are sized to hold approximately three to five days of video archives. Because the failover volumes are contained in the first four archive volume groups, these volumes are necessarily smaller than the archive volumes on archive volume groups numbered 5 through 10. To illustrate this, volume group 1 and volume group 5 are shown in this example:

```
| |- Volume Group VG_ARCHIVE_1 (RAID 6) (32.742 TB)
|   |- Volume VOL_ARCHIVE_1 (28.000 TB)
|   |- Volume VOL_FAILOVER_1 (4,856.000 GB)

| |- Volume Group VG_ARCHIVE_5 (RAID 6) (32.742 TB)
|   |- Volume VOL_ARCHIVE_5 (32.742 TB)
```

Because of this accommodation for the failover volumes, the physical-security integrator should define cameras with lower data rates to the smaller archive volumes if possible.

The Ocularis base machine has a volume (LUN) for bookmarks, which is a central repository for long-term storage of video clips. Given these assumptions, the physical-security integrator has configured the storage array with 26 standard volumes in the following configuration:

```
PROFILE FOR STORAGE ARRAY: stlc2660-33_34 (Thu May 21 15:47:28 EST 2013)

STANDARD VOLUMES-----

Number of standard volumes: 26

Name                Capacity    Accessible by    Source
VOL_ARCHIVE_1       28.000 TB   Host stlc220m3-10 Volume Group VG_ARCHIVE_1
VOL_ARCHIVE_2       28.000 TB   Host stlc220m3-10 Volume Group VG_ARCHIVE_2
VOL_ARCHIVE_3       28.000 TB   Host stlc220m3-10 Volume Group VG_ARCHIVE_3
VOL_ARCHIVE_4       28.000 TB   Host stlc220m3-11 Volume Group VG_ARCHIVE_4
VOL_ARCHIVE_5       32.742 TB   Host stlc220m3-11 Volume Group VG_ARCHIVE_5
VOL_ARCHIVE_6       32.742 TB   Host stlc220m3-11 Volume Group VG_ARCHIVE_6
VOL_ARCHIVE_7       32.742 TB   Host stlc220m3-12 Volume Group VG_ARCHIVE_7
VOL_ARCHIVE_8       32.742 TB   Host stlc220m3-12 Volume Group VG_ARCHIVE_8
VOL_ARCHIVE_9       32.742 TB   Host stlc220m3-12 Volume Group VG_ARCHIVE_9
VOL_ARCHIVE_10      32.742 TB   Host stlc220m3-12 Volume Group VG_ARCHIVE_10
VOL_BOOKMARKS       2,794.000 GB Host stlc220m3-9  Volume Group VG_BOOKMARKS
VOL_FAILOVER_1      4,856.000 GB Host stlc220m3-9  Volume Group VG_ARCHIVE_1
VOL_FAILOVER_2      4,856.000 GB Host stlc220m3-9  Volume Group VG_ARCHIVE_2
VOL_FAILOVER_3      4,856.000 GB Host stlc220m3-10 Volume Group VG_ARCHIVE_3
VOL_FAILOVER_4      4,856.000 GB Host stlc220m3-11 Volume Group VG_ARCHIVE_4
VOL_LIVE_1          1,024.000 GB Host stlc220m3-10 Volume Group VG_LIVE_1_2
VOL_LIVE_2          1,024.000 GB Host stlc220m3-10 Volume Group VG_LIVE_1_2
VOL_LIVE_3          1,024.000 GB Host stlc220m3-10 Volume Group VG_LIVE_3_6
VOL_LIVE_4          1,024.000 GB Host stlc220m3-11 Volume Group VG_LIVE_3_6
VOL_LIVE_5          1,024.000 GB Host stlc220m3-11 Volume Group VG_LIVE_3_6
VOL_LIVE_6          1,024.000 GB Host stlc220m3-11 Volume Group VG_LIVE_3_6
VOL_LIVE_7          1,024.000 GB Host stlc220m3-12 Volume Group VG_LIVE_7_10
VOL_LIVE_8          1,024.000 GB Host stlc220m3-12 Volume Group VG_LIVE_7_10
VOL_LIVE_9          1,024.000 GB Host stlc220m3-12 Volume Group VG_LIVE_7_10
VOL_LIVE_10         1,024.000 GB Host stlc220m3-12 Volume Group VG_LIVE_7_10
```

This configuration has 8 unassigned 3TB drives that can be defined as hot spares. The volume groups utilizing the 900GB 10K RPM drives have 10 drives defined as RAID 10, with a total capacity of 4189.314GB. There are up to four 1TB volumes defined in these volume groups.

During the burn-in phase of the implementation, the physical-security integrator sampled and averaged the data rates from the switch ports supporting the cameras and from the port channel (EtherChannel) to each physical server. The results are shown in the following tabulation:

Physical Server	Host Name	Virtual Machine	Number of Cameras	Portchannel
-----------------	-----------	-----------------	-------------------	-------------

		and Data Rate	Camera	Ingress Rate
Server 4	St1c220m3-12:			275 Mbps
	RACK-SVR-40	64 cameras each 1,383,000 bps	AXIS M3204	
	RACK-SVR-41	64 cameras each 1,383,000 bps	AXIS M3204	
	RACK-SVR-42	64 cameras each 1,678,000 bps	AXIS M3204	
	RACK-SVR-43	64 cameras each 1,678,000 bps	AXIS M3204	
Server 3	St1c220m3-11:			142 Mbps
	RACK-SVR-30	64 cameras each 835,000 bps	AXIS P1346	
	RACK-SVR-31	64 cameras each 1,161,000 bps	AXIS M3204	
	RACK-SVR-32	64 cameras each 1,161,000 bps	AXIS M3204	
	RACK-SVR-33	failover recording server		
Server 2	St1c220m3-10:			225 Mbps
	RACK-SVR-20	64 cameras each 835,000 bps	AXIS P1346	
	RACK-SVR-21	64 cameras each 835,000 bps	AXIS P1346	
	RACK-SVR-22	64 cameras each 1,454,000 bps	AXIS M3204	
	RACK-SVR-23	failover recording server		
Server 1	St1c220m3-9:			
	RACK-SVR-10	OnSSI Ocularis Base		
	RACK-SVR-11	OnSSI Ocularis Manager		
	RACK-SVR-12	failover recording server		
	RACK-SVR-13	failover recording server		

The highest aggregate data rate is on physical server 4, on which all virtual machines have the 32.742TB volumes (LUNs). The 28TB volumes (LUNs) are defined to the three virtual machines on physical server 2 and one virtual machine on physical server 3.

During the implementation phase of the project, as a best practice, camera data rates should be monitored, and cameras should be migrated between recording servers to allocate the load as equally as practical. When the retention period has been reached and the VMS has begun to groom files, the physical-security integrator must verify that the desired retention period is being met or has been exceeded.

4.3 Sizing Example 3: Complex Deployment for a Multiuse Center

In this example, the assumption is that the video surveillance integrator is sizing the storage array to support a multiuse center containing a retail component, self-storage, and offices. The center manager provides physical security services for the complex and, in conjunction with the physical-security integrator, has completed a site survey and determined the number and model of cameras along with the requirement and type of scene.

The physical-security integrator has specified the model and manufacturer of network video cameras and their resolution, frame rate, compression standard and ratio, and image complexity. Axis Communications cameras have been selected for the project. The data rates and storage requirements have been estimated using the Axis Design Tool. The cameras use VBR, and the retention period required is 30 days. The values used to determine the data rates and storage required are shown in Table 4.

Table 4) Multiuse project.

Camera Model	Number	Requirement /Scene	Resolution	Frame Rate	Compression Standard/ Ratio	Image Complexity	Data Rate Mbps	Storage for 30 Days Retention
Axis Q1755	20	Identification cashier	1920x1080	30	H.264/10	Intersection	119.6	36.9TB
Axis Q1755-E	30	Overview parking lot	1920x1080	6	H.264/30	Intersection	30.6	9.4TB

Camera Model	Number	Requirement /Scene	Resolution	Frame Rate	Compression Standard/ Ratio	Image Complexity	Data Rate Mbps	Storage for 30 Days Retention
Axis P3344-VE	50	Identification entrances	1280x720	12	H.264/30	Stairway	52.3	16.1TB
Axis M3204	100	Overview hallways	1280x720	6	H.264/50	Stairway	52.1	16.1TB
Total	200						254.8	78.7TB

Based on the recommendation of the VMS vendor and the experience of the physical-security integrator, four virtual recording servers are deployed on two physical servers. Each physical server will have two failover server virtual machines. In the event one physical machine fails, the failover virtual machines will recover the video streams from the failed server.

In Table 5, the aggregate data rates and storage are shown on a per-camera basis. This data is used to determine the storage required per server.

Table 5) Data rate and storage per camera.

Camera Model	Data Rate per Camera	Storage per Camera for 30 Days Retention
Axis Q1755	5.9Mbps	1843GB
Axis Q1755-E	1.0Mbps	323.2GB
Axis P3344-VE	826Kbps	255.2GB
Axis M3204	372Kbps	114.9GB

The cameras are distributed across the four virtual machines as shown in Table 6. Based on the values in Table 5, the storage requirements of the number and type of camera are calculated and listed in the minimum storage required column. This column includes the assumption of 20% free space per volume (LUN).

Table 6) Camera assignment per server.

Physical Server	Virtual Machine 1 Number/Model Camera	Minimum Storage Required (TB)	Virtual Machine 2 Number/Model Camera	Recorders	Failover 1	Failover 2
One	5 Q1755 7 Q1755-E 13 P3344-VE 25 M3204	21.625TB	5 Q1755 8 Q1755-E 12 P3344-VE 25 M3204	21.625TB	3.6TB	3.6TB
Two	5 Q1755 7 Q1755-E 13 P3344-VE 25 M3204	21.625TB	5 Q1755 8 Q1755-E 12 P3344-VE 25 M3204	21.625TB	3.6TB	3.6TB

Assuming 20% free space, $17.3\text{TB}/0.8 = 21.625\text{TB}$. The failover sizing assumes five days retention based on the 21.625TB value, that is, $21.625 \times (5/30)$ or 3.6TB.

Using Table 2 to estimate the number of disks required to meet the storage needs based on a 60-disk shelf, RAID 6 can be used for the volumes (LUNs) of the active recording servers and RAID 5 for the failover servers and to maintain 4 hot spares. Using a RAID 5 level of protection saves 8 disks compared to using RAID 6 for the failover volumes (LUNs), while still maintaining one volume per volume group. An alternative is to use an 8-disk (6+2) RAID 6 volume group and allocate the four volumes in the volume group. The solution using a combination of RAID 6 and RAID 5 is shown in Table 7.

Table 7) Sizing solution.

Physical Server	VM1 Volume (LUN)	VM2 Volume (LUN)	Failover 1	Failover 2
One	RAID 6 (9+2) = 11 disks 24.556TB	RAID 6 (9+2) = 11 disks 24.556TB	RAID 5 (2+1) = 3 disks 5.5TB	RAID 5 (2+1) = 3 disks 5.5TB
Two	RAID 6 (9+2) = 11 disks 24.556TB	RAID 6 (9+2) = 11 disks 24.556TB	RAID 5 (2+1) = 3 disks 5.5TB	RAID 5 (2+1) = 3 disks 5.5TB

This configuration uses a total of 56 drives, and the four remaining disks can be provisioned as hot spares.

Because this exercise maintains a minimum of 20% free space when selecting the number of drives per volume group, the actual amount of free space per volume (LUN) approaches 30% because it is rounded up to the next number of disk drives.

5 Sizing Checklist

The following checklist describes the data that must be collected and analyzed to determine the capacity requirements of the implementation for accurate solution sizing.

Item	Comments
Retention period	Determines the required minimum retention period of the organization.
Reserve capacity	Estimates how much additional capacity should be allocated in excess of the minimum retention period (for example, retention period plus 20%).
RAID level	Determines what RAID levels are required (for example, RAID 5 or RAID 6).
Disk characteristics	Determines the size and type of disks (for example, if 3TB NL-SAS or 900GB 10K RPM disks are required).
Video management software	Determines which video management system will be implemented.
Volume limitations	Analyzes file system limitations of the video management system or operating system (for example, the Linux 16TB volume limitation).
Data rate per camera	Estimates the data rate expected for each type of camera deployed.
Number of cameras	Determines the total number of cameras and categorizes them by like data rates.

Item	Comments
Continuous recording	Determines if video is recorded continuously or records on motion or on some other analytic trigger.
Frequency of viewing	Determines the frequency, number of concurrent cameras, or video streams being viewed and the number of client-viewing workstations.
Failover requirements	Determines if failover recording servers will be implemented, how many servers, and the size of the volumes required per server.
Cameras per recording server	Estimates the number of servers needed for the specified cameras and the resulting data rates.
Number and size of volumes per recording server	Determines the size and number of volumes per recording server based on the VMS requirements, retention period, and reserve capacity.
Calculation of total disks required	Calculates the number of physical disks required based on the volume sizes per recording server and number of recording servers.
Additional volumes required	Determines if nonrecording volumes are required and what capacity (for example, bookmark volumes).
Hot spares	Determines if sufficient disks are available as hot spares.
Disk shelves required	Determines the number of disk shelves required based on the total number of disks required.

6 Summary

A video surveillance solution based on NetApp E-Series offers the physical-security integrator a highly scalable repository for video management systems.

Sizing the storage array requires knowledge of policy requirements such as the retention period; it also requires knowledge of how many cameras are planned for the deployment and their configuration. Every VMS package has different performance requirements and options on how to configure storage, based on the retention and failover requirements.

Additionally, unknowns exist in each deployment that requires a sizing methodology that allows for some free space in the sizing process, in the event that the estimates for scene complexity are lower than expected. Because of the critical nature of video surveillance to loss prevention and personal safety, high availability must be a key component.

Definitions

This section contains the glossary of terms used throughout this document.

Term	Definition
Storage array	The storage array is a collection of both physical components and logical components for storing data. Physical components include drives, controllers, fans, and power supplies. Logical components include volume groups and volumes. The storage management software manages these components.
Controller	The controller is composed of the hardware board and firmware that manage the physical disk drives and present that capacity to a computer as logical units (LUNs).

LUN	The logical unit number is an address number for how the server identifies different hard drives or, in the case of storage systems, different volumes. Most operating systems will show LUNs as properties of the SCSI hard drives discovered.
Volume group	A volume group is a set of drives that the controller logically groups together to provide one or more volumes to an application host. All of the drives in a volume group must have the same media type and interface type.
SAS	Serial-attached SCSI (SAS) is a computer bus used to move data to and from computer storage devices such as hard drives and tape drives. SAS depends on a point-to-point serial protocol that replaces the parallel SCSI bus technology.
NL-SAS	Near-line-SAS, or NL-SAS, drives are enterprise SATA drives with a SAS interface, head, media, and rotational speed of traditional enterprise-class SATA drives with the fully capable SAS interface typical for classic SAS drives.
RAID	RAID is an acronym for Redundant Array of Independent Disks, and it determines how data is protected from hard drive failures.
RAID 10	RAID 10 provides high availability by combining features of RAID 0 and RAID 1. RAID 0 increases performance by striping volume data across numerous disk drives. RAID 1 provides disk mirroring, which duplicates data between two disk drives. By combining the features of RAID 0 and RAID 1, RAID 10 provides a second optimization for fault tolerance.
RAID 5	A striped disk with parity, RAID 5 combines three or more disks in a way that protects data against loss of any one disk. The protected storage capacity of the volume group is reduced by one disk from the raw capacity.
RAID 6	Striped disks with dual parity, RAID 6 can recover from the loss of up to two disks. The protected storage capacity of the volume group is reduced by two disks from the raw capacity.
HDTV	High-definition TV defines resolutions of 1920x1080 and 1280x720 pixels along with other criteria, including aspect ratio.
Megapixel	This is any video resolution of 1 million pixels or more. However, the HDTV resolution of 1280x720 is 921,600 pixels but is commonly referred to as a megapixel resolution.
M-JPEG	A video codec—Motion Joint Photographic Experts Group.
MPEG-4	A video codec—ISO/IEC Moving Picture Experts Group (Part 2)—and the predecessor to H.264.
H.264	A video codec—ISO/IEC Moving Picture Experts Group (Part 10)—and is more efficient than MPEG-4 and commonly used by HDTV and megapixel cameras.
VMS server	Video Management System Server is also referred to as Network DVR Recording Server. It manages IP camera video feeds and storage media.
Viewing station	A high-end workstation for displaying live or archived camera feeds on a locally attached monitor.
IP video surveillance camera	A digital video camera or network video camera is a small form factor IP networked Linux host that encodes and transports video over an IP network.

References

The following references were used in this document:

- NetApp University Customer Learning Map
http://learningcenter.netapp.com/content/public/production/learning_maps/customer/lm_customer_t2.html
- Video Surveillance Storage Solution
<https://fieldportal.netapp.com/video-storage.aspx#14791>

- Axis Communications
<http://www.axis.com/>

Version History

Version	Date	Document Version History
Version 1.0	July 2013	Initial release

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®



www.netapp.com

© 2013 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, and SANtricity are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Linux is a registered trademark of Linus Torvalds. Microsoft, Windows, and Windows Server are registered trademarks of Microsoft Corporation. VMware is a registered trademark and ESXi is a trademark of VMware, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4199-0713