



Milestone Systems – White Paper

Storage Architecture in XProtect[®] Corporate

Prepared by: John Rasmussen – Product Manager XProtect Corporate

Milestone Systems

Date: 7 February, 2011

Table of Contents

| | |
|--|----|
| Introduction | 3 |
| Target audience and purpose | 3 |
| Managing high volume of video data efficiently | 4 |
| The importance of managing high volume of video data efficiently | 4 |
| The specific nature of video surveillance data..... | 4 |
| The ability to utilize low cost storage technologies | 5 |
| Heavy duty..... | 5 |
| Technical overview..... | 6 |
| Introduction..... | 6 |
| Storage architecture in XProtect Corporate 3.x..... | 6 |
| Storage architecture and storage features in XProtect Corporate 4.0..... | 7 |
| 1. Multi-stage archiving | 7 |
| 2. Encryption | 9 |
| 3. Video grooming..... | 10 |
| 4. Performance..... | 11 |
| 5. Background repair..... | 14 |
| 6. Storage dashboard | 15 |
| Benefits of XProtect Corporate’s storage architecture..... | 16 |
| Storage Configurations and Recommendations | 16 |
| Drive types..... | 16 |
| RAID configuration | 17 |
| Archiving vs. no archiving..... | 17 |
| Upgrading to XProtect Corporate 4.0 | 18 |
| Summary | 19 |

Introduction

This white paper describes the storage architecture in Milestone XProtect Corporate IP video management software, and the new additions introduced in release 4.0 of XProtect Corporate.

Milestone XProtect Corporate supports a unique storage architecture for video surveillance systems, which offers the system administrator or integrator the flexibility to design the storage system to meet the customers' demands for performance, reliability and price.

Previous releases of XProtect Corporate supported a two-stage storage solution, where audio and video could be recorded to one storage system and optionally archived to another storage system after a pre-defined time and schedule.

XProtect Corporate 4.0 extends the existing storage architecture with support for multi-stage archiving, database encryption and video grooming.

In addition to this, the surveillance database has been refined to significantly reduce database file fragmentation on the recording storage system: the database now will write each database file in as few fragments as possible so fragmentation has been reduced from typically 2-300 fragments per file to just a few (typically one). This hugely improves the performance of the storage system and ensures that the disk performance is kept constant over time as the disk is rewritten when recordings are archived or deleted.

Target audience and purpose

The target audience of this white paper is surveillance system architects/designers, surveillance project consultants, companies and governments with surveillance installations.

The purpose of this white paper is to give a general overview of Milestone XProtect Corporate's storage architecture, insights and recommendations to design and configuration of XProtect Corporate's surveillance database and storage system.

This white paper should enable the reader to understand the architecture and technology behind the surveillance database and archives in XProtect Corporate, as well as understanding how to design and implement a storage system for an XProtect Corporate surveillance system.

The white paper assumes that the reader has a general understanding of Milestone XProtect Corporate, its surveillance database, archives, and different storage technologies and solutions.

Managing high volume of video data efficiently

To put Milestone's unique storage and archiving implementation in context, and to better understand the importance of efficient video storage management, let's start by taking a closer look at the role of video storage as a part of the greater surveillance solution. Doing so, we'll also look at how the specific characteristics of video data make it quite different from other data known from IT systems.

The importance of managing high-volume video data efficiently

In larger surveillance installations the storage infrastructure normally accounts for 10-15% of the overall system cost, dependent on the specific conditions and retention time of the video. This means that relatively large savings can be made with even modest degrees of improvement in the storage handling and storage platforms.

To illustrate this, picture a surveillance installation with 700 cameras, where video recordings are kept 3 months for reference purposes due to legal requirements. However, possible incidents are normally investigated within 14 days, and never later than a month after the occurrence of the incident.

Such a system could easily generate 6 Terabyte of video recordings per day¹, which would correspond to 550 Terabyte of video data for the full 3-month period in our example. However, if we optimize the data during the last 2 months when the video is kept for reference purposes only, significant savings can be made. Let's say that it may be enough to maintain the video with 1 Frame Per Second (FPS) after the initial month, which can reduce the data volume and hence the overall storage cost by over 50%.

The specific nature of video surveillance data

On first consideration it is easy to make the mistake of considering video data as any other IT or business data, where the only difference is the volume of data. However, there are two aspects that differentiate video surveillance data from other business data, namely:

- Video data is streamed real-time data
- Multiple streams are handled and stored in parallel

In a standard IT system, the disks are used to store files of data that are read from or written to by other servers, IT systems or users. When files are written to these disks, it is done using "best effort writing", and it does not really matter if the disks are fast enough, as the data also is stored or buffered by the sender - the write just takes a little longer, but data is never lost. Also, the files to a large extent are written sequentially one by one or just a few in parallel when more systems or users access the disk system at the same time.

¹ Based on the assumption that 200 cameras run HD H.264, and 500 cameras run 4CIF JPEG with 5 frames per second, and with an average motion/recording rate of 15%.

In a surveillance system, video is received as continuous real-time data streams, from many parallel sources (cameras), and there is nowhere data temporarily can be stored or buffered if the disk system is too slow or busy. This means video recordings are lost if the disk system cannot keep up all the time.

The disk load will also be massive as the disk needs to write the incoming data in parallel to many different database files at the same time. This parallel writing is called non-sequential writing and causes most standard disks to perform slower compared to sequential writing where only one file is written at a time.

The ability to utilize low-cost storage technologies

A number of storage technologies with different performance abilities exist today, including SATA (Serial Advanced Technology Attachment), SCSI (Small Computer System Interface) and SAS (Serial Attached SCSI). Common to these technologies is that the cost of the actual hard disk drives (HDD) are proportional to the performance (normally measured in Input/Output operations Per Second - IOPS), while the cost per Gigabyte varies significantly between the lower specification SATA HDDs and the higher performance SAS and SCSI HDDs. The difference in price per Gigabyte for a SATA HDD compared with a SAS or SCSI HDD has been 1 to 4 for quite some time.

To manage the high volumes of parallel writings that are characteristic for surveillance video disc technologies with high Input/Output operations Per Second (IOPS), SAS or SCSI are normally required. However, a significant portion of the storage cost can be saved if the video surveillance system can utilize a storage infrastructure primarily based on SATA HDD. Using SATA-based storage exclusively, instead of SAS or SCSI disk technologies, can save as much as 75% on storage costs!

Heavy duty

When determining the total cost of ownership for the storage infrastructure it is important to consider the considerable mechanical wear on the drives caused by the continuous 24/7/365 writing of data. This means that the investment cost should be considered together with the expected lifetime of the HDD used, and its replacement cost. Typically SAS/SCSI HDD are more robust, and therefore have longer expected life.

Technical overview

This section describes the storage and archiving architecture in XProtect Corporate 3.x and earlier versions of the product, and the new storage management concepts and features introduced in XProtect Corporate 4.0

Introduction

The storage architecture consists of a proprietary Milestone surveillance database specifically designed and optimized for video surveillance usage with efficient storage of multiple real-time video streams with optional archive functions to seamlessly move the surveillance databases to online archive drives.

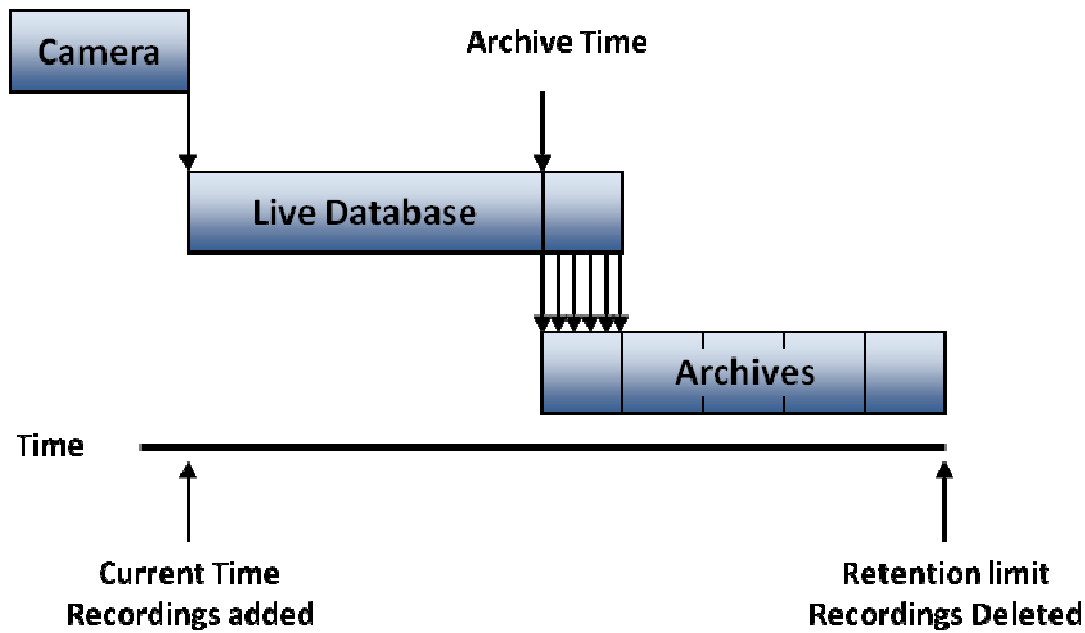
The architecture with the archiving function has been developed to ensure the best possible performance and price for the storage system, even in very large video system installations with thousands of cameras and storage needs ranging from Terabytes to Petabytes.

Storage architecture in XProtect Corporate 3.x

In XProtect Corporate release 3.x, and earlier versions, the Recording Server database supported a two-stage recording/archiving solution where audio & video recordings first are stored in the so called 'Live' database, and then optionally after a set time and schedule are moved to an online 'Archive' database.

With this architecture the recordings are always online for the users of the system, who can seamlessly access the recordings from either the Live database or the Archives - even while archiving is occurring.

The drawing below illustrates how video and audio recordings for one camera are retrieved from the camera, stored in the Live database, archived based on a schedule, and finally deleted when the retention times are reached.



Storage architecture and storage features in XProtect Corporate 4.0

XProtect Corporate 4.0 introduces a number of new features to the storage architecture known from the previous XProtect Corporate versions.

1. Multi-stage archiving

Where older versions (version 3.1 and earlier) of XProtect Corporate support the two-stage storage architecture, XProtect Corporate 4.0 now supports a multi-stage storage architecture where recordings can be archived again and again to new storage areas.

The multi-stage archiving combined with the new grooming feature (see below) measurably reduces the storage costs for systems that require recordings to be stored for a long time. This is enabled by a gradual reduction of the video frame rate and hence the data volume over time.

In addition to the multi-stage archiving and grooming features, the 'Live' databases are now automatically divided into 1-hour segments. This minimizes the size of the open 'Live' databases in order to manage a potential database repair after a failure as quickly as possible. See also background repair section below.

The definition and configuration of 'Live' databases and archives now are defined as part of a new concept called a 'Storage container'. Cameras and devices that should store video or audio recordings are set to use one of the defined Storage containers, making the storage configuration on the individual devices very simple.

The drawing below illustrates a Storage container with a live database and 3 archives.

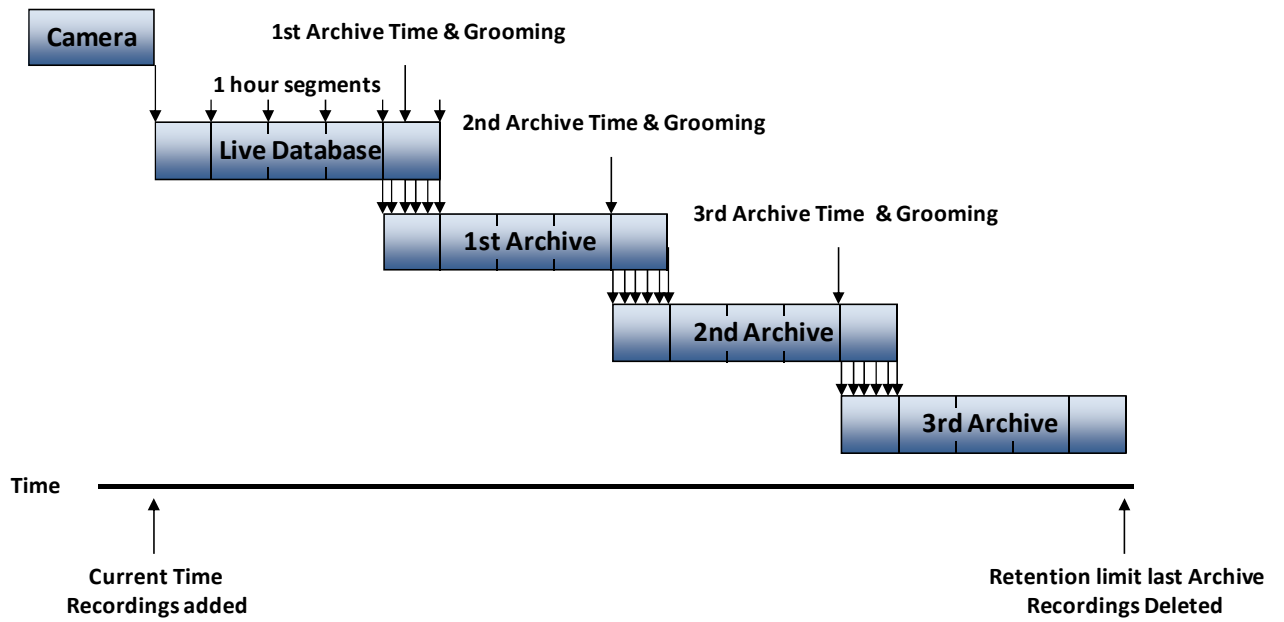


Figure 2: XProtect Corporate 4.0 multi-stage storage management

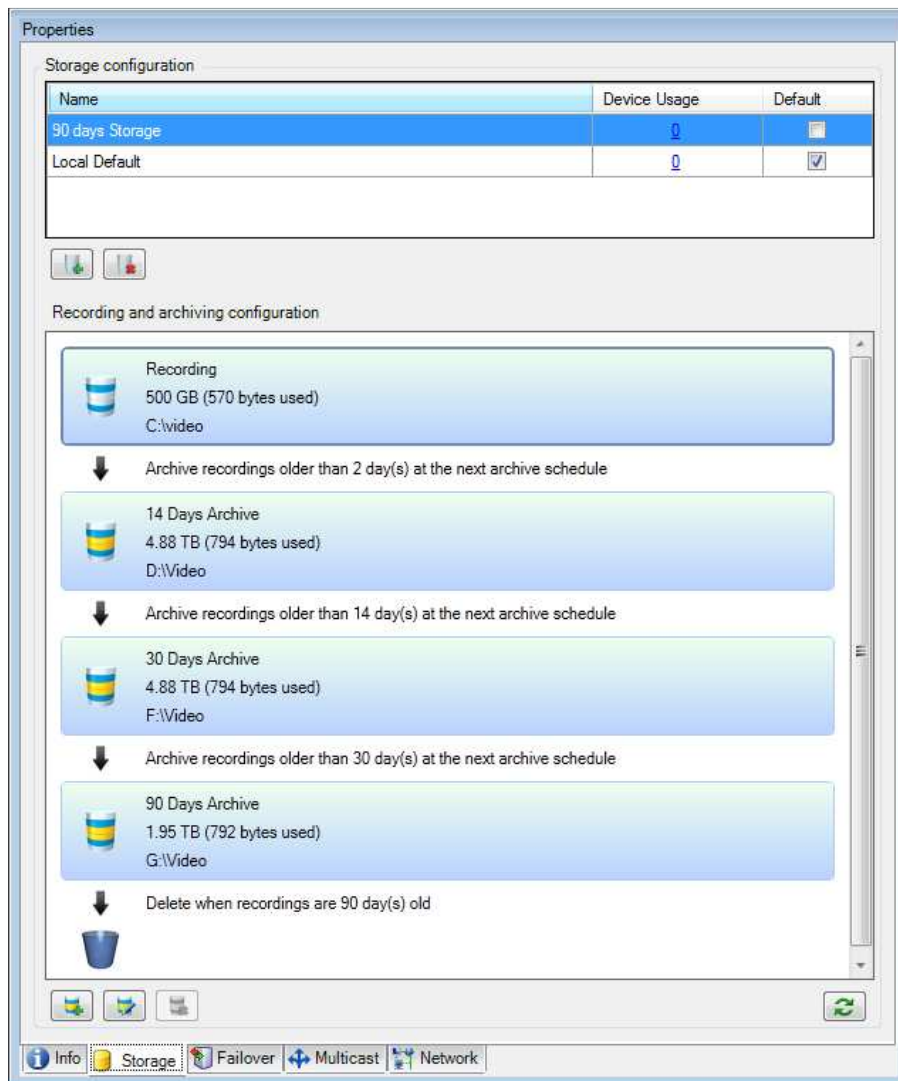


Figure 3: Screenshot of the Management Client showing storage configuration with a Live database and three Archives

2. Encryption

XProtect Corporate 4.0 also introduces encryption of video and audio recorded in the surveillance databases.

Encryption can be enabled per Storage container, which then encrypts video and audio data recorded in the live database from the cameras assigned to this Storage container.

The encryption is kept and transferred with the recordings once they are archived. If the archiving function also grooms the video recordings, the encryption is still kept since it is the individual records (images, Group of Pictures - GOP's or audio) inside the database that are encrypted.

Encryption is done using a Data Encryption Standard (DES) 56 bit block cipher, and can be enabled as either full encryption of the database records or just as encryption of

the record headers. Encrypting just the headers is enough to prevent unauthorized decoding of the video or audio since it cannot be decoded without the information in the header.

Enabling encryption adds extra CPU load to the Recording Servers as recorded audio and video needs to be encrypted. Choosing to encrypt only the headers is cheaper CPU-wise compared to encrypting the full record.

The graph below shows the CPU load on a system running 32 cameras @ 25 FPS & 47KB MJPEG images, recording all images.

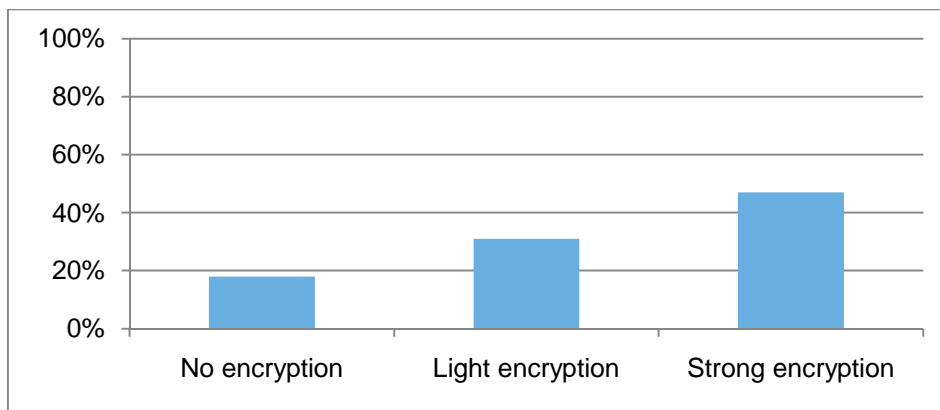


Figure 4: Relative CPU load for the use of light versus strong encryption²

3. Video grooming

XProtect Corporate 4.0 introduces the possibility to groom the video recorded in the surveillance databases each time the recordings are archived.

Grooming is a method to decrease the frame rate of the recorded video over time in order to save space in the storage system while still keeping a record of what has happened. The typical argument for using grooming is: the older the video, the less likely it is that something important is recorded on it because important things are typically discovered and investigated quickly.

XProtect Corporate 4.0 can groom video recordings each time they are archived and moved from one storage area to another. Grooming is enabled with the specific frame rate setup in the same dialog where archiving is configured. Since XProtect Corporate 4.0 supports multi-stage archives, the grooming can be done as many times as there are archives, reducing the frame rate again and again over time.

For JPEG recordings it is possible to groom to any frame rate lower than the frame rate currently recorded in the database. For MPEG and H.264 recordings, grooming can only be done to key-frames and lower (e.g. a key-frame every 1 second or less).

² Server used for measuring the encryption CPU load was HP xw6600 with Intel XEON E5430 @ 2.66 GHz (quad core), 4GB RAM, Windows Server 2008 R2 - 64 bit, SATA disk - Samsung HD502HJ 7200 rpm.

4. Disk Performance

A disk system with conventional rotating disk technology will perform slower the more fragmented the files on the disk are. Hence, a key parameter to ensure continuous high disk performance is to keep the file fragmentation as low as possible.

XProtect Corporate's live database consists of multiple data block files (*.blk) per hardware device (camera) that contain the individual recordings. In previous versions of XProtect Corporate, these database files become fragmented over time because Windows dynamically allocates space for them as the system writes video or audio records in the files.

XProtect Corporate 4.0 introduces an optimization to the surveillance database that ensures the database files are stored on the disk system in as few fragments as possible - typically one fragment. This is done without first buffering the video in memory or having to pre-format the disk for surveillance usage, which would lead to increased memory utilization, or long initialization or format time.

The graph below shows the increase in recording performance from XProtect Corporate 3.x to 4.0 when recording to a single SAS or SATA disk.

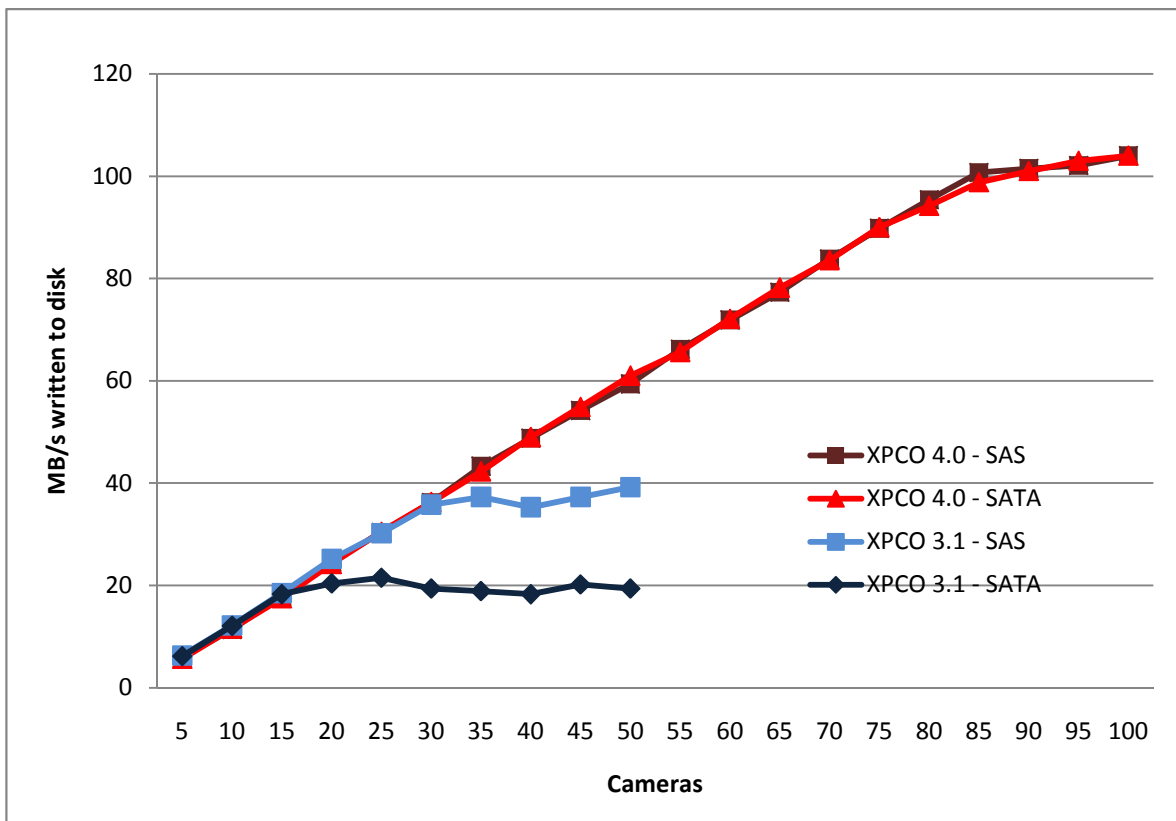


Figure 5: Recording Performance Improvement in XProtect Corporate 3.1 vs. 4.0

For this test the disks were formatted with standard 4K file allocation size and cameras were configured to MJPEG 4CIF 47 KB images @ 25 FPS - recording all images.

The graph shows that with XProtect Corporate 3.x the disks become saturated earlier due to increased disk seeking caused by file fragmentation - whereas XProtect Corporate 4.0 utilizes the disks much better because each file is stored in one single fragment, which reduces the disk seeking time when there are many cameras recording to the same drive.

Disk idle time

Another way of measuring the disk performance is to monitor the disk's idle time when performing a given set of write or reading tasks. The more fragmented the disk becomes the longer it takes to write and read data to/from it, and the shorter the idle time becomes. By measuring the disk idle time over a defined period, it's possible to measure the sustained performance. If the performance degrades over time, this is due to increased file fragmentation.

The graph below shows the idle time of the disk used for the live database in XProtect Corporate 3.x compared to 4.0. In both cases the system was set up to record video data corresponding to 40 MB/s being written to the disk, and with hourly archiving.



Figure 6: Live disk idle time measured over time for XProtect Corporate 3.x and 4.0³

As shown by the graph above, XProtect Corporate 4.0 has a higher idle time compared to XProtect Corporate 3.x. The difference indicates that XProtect Corporate 4.0 utilizes the disk approximately 45% better than XProtect Corporate 3.x.

It can also be seen that the disk idle time stays constant for XProtect Corporate 4.0, whereas it drops somewhat for XProtect Corporate 3.x. This indicates that the file

³ Server used for measuring the disk performance was a HP xw6600 with: Intel XEON E5430 @ 2.66 GHz (quad core), 4GB Ram. Windows Server 2008 R2 - 64 bit, SATA disk - Samsung HD502HJ 7200 rpm and SAS disk - Seagate Cheetah 15K6.

fragmentation does not increase over time with XProtect Corporate 4.0, thus ensuring that the high disk performance is sustained.

An additional benefit of the increased performance in XProtect Corporate 4.0 is derived from the fact that the higher a disk's idle time is, the lower the mechanical wear, and the longer the disk will last. This translates to fewer disk failures and lower total cost of ownership.

Disk formatting

Another factor that can impact a disk's performance in a video surveillance system is the cluster size of the disks' formatting, where 64K file allocation size gives better performance than disks formatted with the standard 4K.

The below graph shows the recording performance to a Seagate Cheetah SAS disk formatted with standard 4K and 64K file allocation.

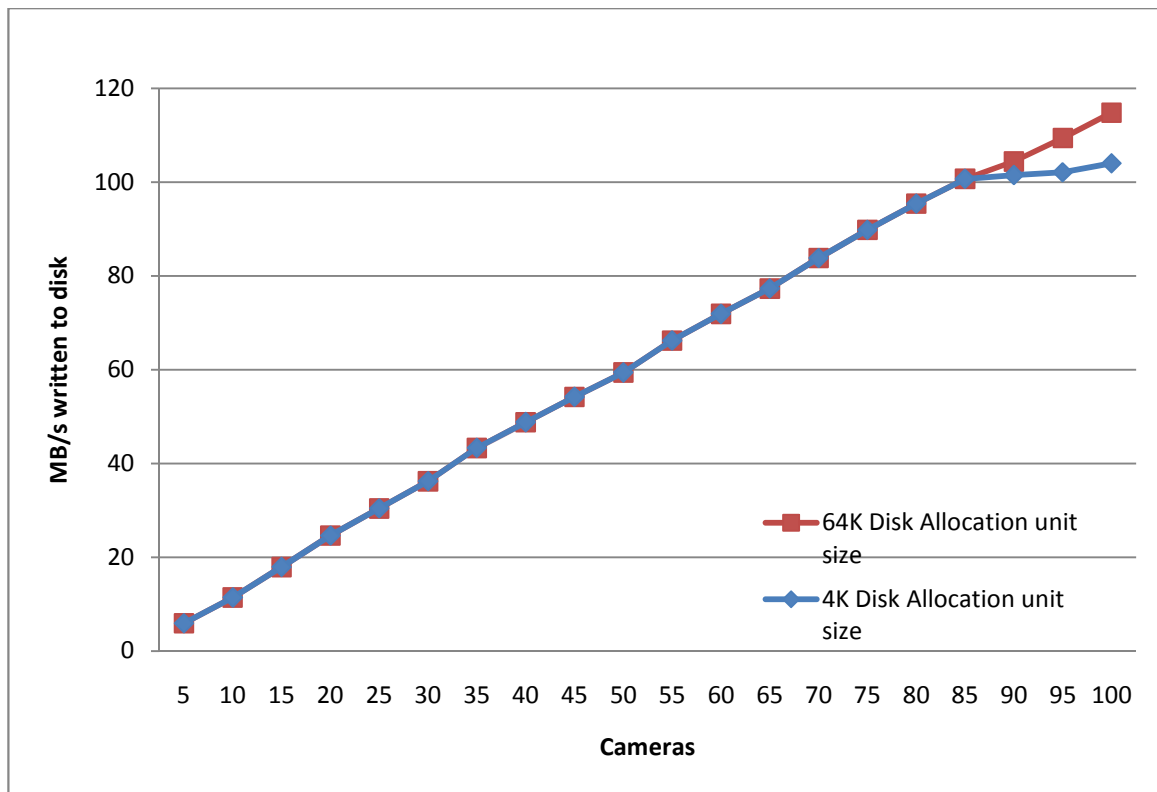


Figure 7: Recording Performance 4K vs. 64K formatting Seagate Cheetah 15K rpm SAS disk⁴

As indicated in the graph above, the disk performance with 64K formatting exceeds the performance with standard 4K formatting when connecting a high number of cameras to the same Recording Server. Not shown in this diagram is that the disk idle time also

⁴ Server used for measuring the disk performance was a HP xw6600 with: Intel XEON E5430 @ 2.66 GHz (quad core), 4GB Ram. Windows Server 2008 R2 - 64 bit, SATA disk - Samsung HD502HJ 7200 rpm and SAS disk - Seagate Cheetah 15K6.

increases with the 64K formatting, giving a better responding disk, even for systems that do not load the disk to the same extreme as shown in the above graph.

It is therefore recommended to format the disks used in XProtect Corporate systems with a 64K file allocation size because it increases both performance and idle time.

5. Background repair

XProtect Corporate 4.0 introduces the ability for the Recording Servers to repair the databases in the background, both in start-up scenarios and on the fly, if it detects databases that might be corrupted.

In start-up scenarios where the Recording Server was not shut down properly, it could previously take a long time to start the Recording Server (depending on size and condition of the database) since the databases needed to be repaired first.

The start-up procedure has now been changed: during a start-up with corrupted databases, they are now moved to a sub-folder, new databases are created, and the Recording Server starts normally. Once the Recording Server is up and running, these corrupted databases are repaired in the background and merged into the new database.

During the database repair process it is indicated in the Management Client for relevant devices and Recording Servers that databases are being repaired.

Users of the Smart Client will experience corrupt databases that are being repaired as gaps in the recordings. Once the databases are repaired, one by one, the recordings they contain can be browsed by the clients without any further actions.

This changed start-up scheme with background repairs ensures that the start-up times for Recording Servers are the same whether there are corrupted databases to be repaired or not.

6. Storage dashboard

XProtect Corporate 4.0 introduces a new Storage Dashboard in the Management Client that gives an excellent overview of the storage usage and current database tasks in the surveillance installation. The storage usage can be viewed both from a Recording Server perspective, and for individual devices.

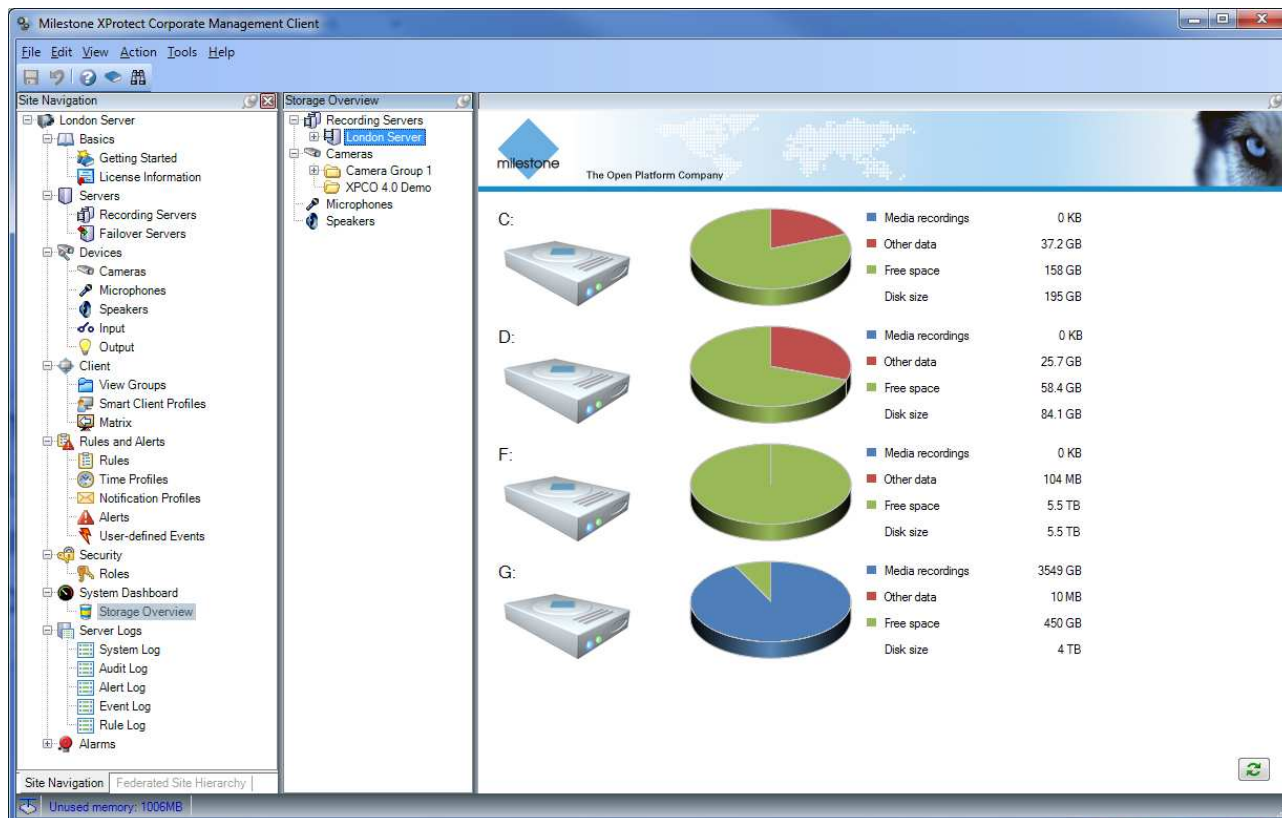


Figure 8: The Storage Overview in the Management Client gives instant overview of the use of available storage

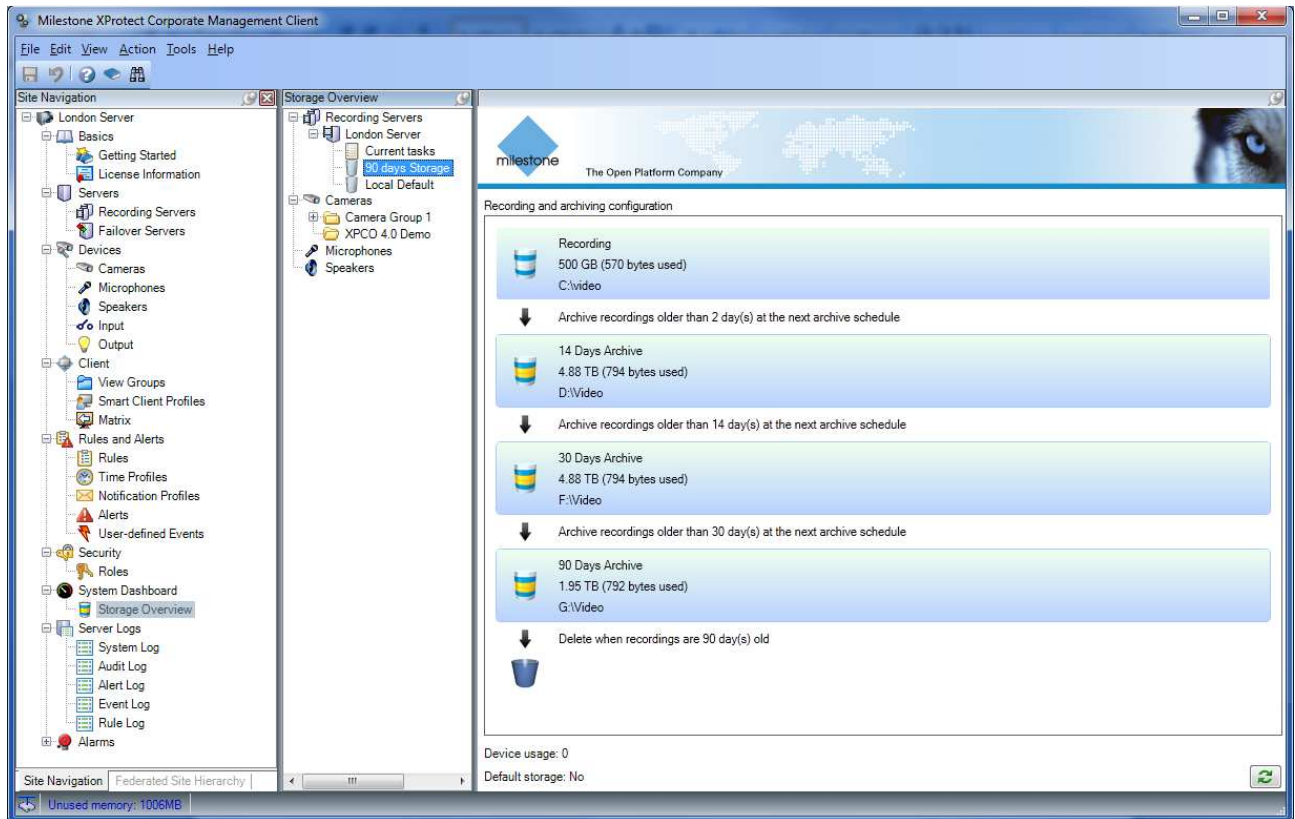


Figure 9: The storage overview also gives instant access to defined storage containers

Benefits of XProtect Corporate storage architecture

The benefit of XProtect Corporate's storage architecture is that it gives surveillance system architects, installers and administrators the flexibility to design the storage system to their specific needs - while getting the best of both worlds: high performance and endurance on the 'Live' disks with low cost and large capacity for archives.

The storage architecture also offers the ability to take advantage of the newest storage technologies like Solid-State Drive (SSD), which although still somewhat expensive are ideal for 'Live' storage because their performance is optimal even for non-sequential writing and they do not wear out mechanically.

Storage Configurations and Recommendations

Drive types

The storage recommendation for a typical Milestone XProtect Corporate surveillance system with archiving is to use SCSI or SAS disk technology for the 'Live' disks because they offer a good non-sequential write performance and mechanical endurance, and to use larger, relatively inexpensive server class SATA disks for the Archive disks.

Optionally, the live disks can be SSD disks which do not have any moving parts, and thus do not wear out mechanically. One common worry with SSD disks are write endurance, but if high endurance SSD disks with wear leveling are used, there should be no trouble in getting the disks to last 10 years or more.

Once the recordings are stored on the live disks, they can be archived one database at a time to the archive disks. This way provides better disk performance as the archive writings now are sequential (one camera at a time), and if the disks are busy for a second or two, the data is not lost because it is stored in the live database until all of it has been archived. This also makes it possible to use an Redundant Array of Independent Disks (RAID) 5 or 6 configurations for the archive drives.

RAID configuration

Typically the disks, both Live and Archive, are configured to use RAID in order for them to be large enough. In this case it is recommended to only use RAID 0, 1 or 1+0 for the 'Live' disks and then RAID 5 or 6 for the archive disks.

The reason for only recommending RAID 0, 1 & 1+0 for the live disks is that there is a large overhead on the RAID controller to calculate the parity and writing the data at a high rate in parallel across all the disks in a RAID 5 or 6 disk system.

Normally it is not a problem to use RAID 5 or 6 in a normal IT environment, as buffers in both ends hold the data until it is written to the disk. However, in a video surveillance system there are no buffers because the data (video) is received in real-time. This means that the disks and controllers continuously need to be able to handle the data in real-time, which often proves to be more than most controllers can handle in RAID 5 or 6, even on more expensive disk systems.

The recommendation for the drive type usage and RAID configuration is also the same when the drives are used in a storage box like a Storage Area Network (SAN), Direct-Attached Storage (DAS), Network-Attached Storage (NAS), etc.

Archiving vs. no archiving

XProtect Corporate offers the option to not archive video recordings, and simply let them reside in a single, very large database for each camera.

In the previous XProtect Corporate 3.x release this was not recommended because a database repair after a failure could potentially take a very long time, and would have to be completed before the Recording Server could start and resume recording.

In XProtect Corporate 4.0 this disadvantage has been eliminated with the live database automatically divided into 1-hour segments. This reduces the data that potentially should be repaired to no more than the last hour's segment, no matter what size the 'Live' database overall may have. In addition, the Recording Server can start even with a corrupt database and repair it in the background ensuring a normal startup time.

These improvements to the database and repair process means that it does not matter if the databases are archived or not, so it is up to the system integrator or administrator to choose if archiving should be used or not. If archiving is not used, the disks need to be fast and large enough to record and store the video for the full retention period required.

As the decision whether or not to use archive depends on a number of factors such as: data volumes, used hardware and disk technology, the use of data grooming, etc. it is not possible to give any guidelines on when to use archives as compliments to the Live database.

Upgrading to XProtect Corporate 4.0

When upgrading to XProtect Corporate 4.0 it is important that the whole system is upgraded altogether at once, as there are some fundamental changes to both the Management Server and Recording/Failover servers, which cause the components in the previous versions to be incompatible with the 4.0 components.

When upgrading the Management Server to 4.0 the old database and archive configuration for the Recording Servers and cameras will be imported and converted to the new Storage container configuration. A Storage container will be created for each permutation of database/archive for the different cameras on each Recording Server. In a typical installation there will only be a few database/archive configurations for the cameras on each Recording Server, so there should also only be a few Storage containers after the upgrade.

When the Recording Servers afterwards are upgraded to the 4.0 version, the existing recordings and archives for the cameras will be kept and continue to be accessible.

Summary

The ability to store high volumes of IP video surveillance data efficiently is not only critical for the overall performance and reliability of the surveillance solution, but physical storage often accounts for a significant part of the total system cost - a cost that can be reduced dramatically with the right tools.

In this paper we've illustrated how XProtect Corporate 4.0's new multi-stage storage management and data grooming can be used to drastically cut the cost of long-term video storage, by reducing the frame rate of video stored over longer periods.

XProtect Corporate 4.0 also offers remarkable performance improvements on the actual data writing process, which means that more high resolution video data can be written to the database per second, compared to earlier versions. Disk performance is sustained over time with disk fragmentation kept to a bare minimum.

The performance improvements also translate to reductions in the total cost of ownership, because the enhanced database is less sensitive to the actual storage technology used. In most installations XProtect Corporate 4.0 will perform fully satisfactorily on lower specification SATA storage infrastructure. In normal operation the improved data writing performance creates less stress on the individual disk drives so they can last longer.

The final key aspect of the new multi-stage storage in XProtect Corporate 4.0 is the enhanced recovery handling, where the recovery of a corrupted live database in the background makes the Recording Server able to record new video almost instantly after the detection of the faulty database.



The Open Platform Company

About Milestone Systems

Founded in 1998, Milestone Systems is the global industry leader in true open platform IP video management software. The XProtect™ platform delivers powerful surveillance that is easy to manage, reliable and proven in thousands of customer installations around the world. With support for the widest choice in network hardware and integration with other systems, XProtect provides best-of-breed solutions to 'video enable' organizations – managing risks, protecting people and assets, optimizing processes, and reducing costs. Milestone software is sold through authorized and certified partners. For more information please visit www.milestonesys.com

Milestone Systems Headquarters, DK

Tel: +45 88 300 300

Milestone Systems US

Tel: +1 503 350 1100