

White paper

XProtect Storage Architecture and Recommendations

Prepared by:

John Rasmussen, Platform Architect

Milestone Systems

Date: January 31, 2019

Table of Contents

List of abbreviations and terms.....	4
Introduction.....	6
Purpose and target audience	6
XProtect VMS data vs. IT data.....	7
Standard IT storage system usage.....	7
XProtect VMS storage system usage	8
Summary.....	8
Storage architecture	9
Storage and Media Databases	9
Pre-buffer	10
Recording database.....	11
Archive database.....	11
Storage technologies	12
Standard disk technologies and key characteristics	13
Hard Disk Drive.....	13
Solid-State Drive	14
RAID - Redundant Array of Independent Disks	15
RAID 0.....	16
RAID 1	16
RAID 5.....	17
RAID 6.....	19
RAID 10 (also called RAID 1+0 or RAID 1&0).....	20
Other RAID configurations	22
RAID – hot spare.....	22
RAID controllers	22
Storage enclosure technologies	22
DAS.....	23
NAS.....	23
SAN.....	24
Storage configuration	24
Basic configuration	24
Archiving.....	26
Reduce framerate.....	27
MJPEG video	27

MPEG video	27
Configuration	28
Differentiated retention time	29
Low disk space handling	29
System monitor events	30
Securing XProtect VMS media databases	30
Permissions	30
Media database encryption	31
Media database signing.....	32
Evidence Lock	32
Using Evidence Lock.....	32
Storage impact.....	34
Media database performance and recommendations	36
Recording database.....	36
Disk recommendations.....	37
Recording video only	37
Recording video, audio and metadata	38
Archiving.....	39
Reducing framerate when archiving	39
Database repair.....	40
Multiple storage configurations.....	40
Number of recording servers.....	41
Codec	41
GOP length	42
Windows recommendations	42
Separate OS and XProtect VMS drives.....	42
Disk formatting	43
HDD short stroking	43
Windows search indexing	43
Windows disk defragmenter	44
SSD trim.....	44
Virus scanner	44
Windows update and maintenance.....	44
Troubleshooting.....	45
Summary	46

List of abbreviations and terms

B-frame	Bi-directional frame – contains information about changes to the frame compared to the previous and next frame in a GOP
Codec	Short for 'coder-decoder' – a codec contains information about how to encode and decode audio/video according to the codec's standard – for instance; MPEG-4, H.264 and H.265
CPU	Central Processing Unit - the component in a computer that carries out the instructions of a computer program
Device	Common denomination of cameras, microphones, speakers and metadata sources in an XProtect VMS
FPS	Frames Per Second – number of images per second in the video stream
GOP	Group Of Pictures – collection of successive pictures within a MPEG encoded video stream
GPU	Graphics Processing Unit - the component in a computer that handles graphics operations, like decoding video and rendering graphics
H.264	MPEG video compression standard – successor to MPEG-4 and predecessor to H.265
H.265	MPEG video compression standard – successor to H.264
HDD	Hard Disk Drive – a disk drive with spinning plates, storing the data as magnetic information on the plates
I-frame	Also known as keyframe – the first image in a GOP. An I-frame is comparable to a JPEG image
IOPS	Input/output operations Per Second – a performance measurement used to characterize storage devices
JBOD	Just a Bunch Of Disks - multiple independent hard disk drives
JPEG	Joint Photographic Experts Group – a compression standard for still images
Keyframe	Also known as I-frame – the first image in a GOP. A keyframe is comparable to a JPEG image
Media Data	Video, Audio and metadata retrieved from connected devices
MJPEG	Motion JPEG – a series of JPEG images forming a video stream
MPEG	Moving Picture Experts Group – used as a common term covering MPEG-4, H.264 & H.265 codecs in this whitepaper
MPEG-4	MPEG video compression standard – predecessor to H.264

NAS	Network-attached storage - a storage technology providing access to file level storage
OS	Operating System
PCB	Printed Circuit Board – a board that holds and connects components in electronic devices
P-frame	Predictive frame – contains information about changes to the frame compared to the previous frame in the GOP
RAID	Redundant Array of Independent Disks – a technology that combines multiple physical disks into one or more logical disks, with support for fault tolerance
RPM	Revolutions Per Minute
SAN	Storage Area Network – a storage technology providing access to consolidated, block level data storage
SAS	Serial Attached SCSI – communication standard for storage devices
SATA	Serial ATA – communication standard for storage devices
SSD	Solid-State Drive – a disk drive without any moving parts, storing the data in flash memory chips
VMD	Video Motion Detection – the process of detecting movement in video images
VMS	Video Management System or Video Management Software

Introduction

In a VMS a key function is the ability to efficiently, reliably and securely record and store real-time video, audio and metadata from the connected cameras, microphones, speakers and metadata sources.

The Milestone XProtect VMS offers the world's best performing recording server capable of recording at least up to 3.1 Gigabit/s from several hundred cameras.

The recording server's media database supports all standard disk and storage technologies known from the IT industry and provide the freedom to design and use more or less any type of storage system and architecture ranging from recording to a single SATA drive in a PC to recording to large redundant storage systems with multiple archives and usage of network storage.

Furthermore, the recording server's media database offers secure handling of the recorded media by supporting encryption and digital signing of the media stored in the database. Furthermore, it support protecting recorded media for periods extending the standard retention time by using Evidence Lock.

Throughout this white paper, the products listed below are jointly referred to as "XProtect VMS":

- XProtect Corporate
- XProtect Expert
- XProtect Professional+
- XProtect Express+

Purpose and target audience

The purpose of this white paper is to provide insights on how XProtect VMS data and its usage differs from standard IT data, and how this impacts the storage design and choice of disk and storage technologies.

This white paper should enable the reader to understand:

- The difference between standard IT data and real-time XProtect VMS data
- Disk and storage system technologies and how they work in an XProtect VMS environment
- The XProtect VMS storage architecture – including;
 - Memory based pre-buffer
 - Recording and archive databases
 - Reducing framerate during archiving
 - Securing the XProtect VMS media databases
 - Using Evidence Lock and the impact this has on storage
 - Monitoring system performance

The white paper assumes the reader has a general understanding of Milestone XProtect VMS and standard IT storage technologies.

The primary audience for this white paper might include (but is not limited to) the following audiences:

- Surveillance system architects/designers
- Surveillance project consultants
- Companies, organizations and governments with surveillance projects/installations
- Support technicians and engineers troubleshooting XProtect VMS installations

XProtect VMS data vs. IT data

Before covering the details of the Milestone XProtect VMS media database, storage architecture and storage system technologies, it is important to understand the nature of XProtect VMS media data in the form of video, audio or metadata being streamed, recorded and stored in real-time.

At first glance, without knowing the details of how an XProtect VMS work, an IT system designer or administrator will often think of XProtect VMS media data as being the same as standard IT data like documents or picture files stored on a corporate network share - with the amount of data being the only difference.

If considering a static XProtect VMS that is not recording, this belief is correct. In that case the media database is just files stored in folders.

However, an XProtect VMS is never static. Every single second it writes new media data retrieved from cameras and other devices to the storage system in real-time. Furthermore, media older than the set retention time is also constantly being deleted or archived to a new location.

Because of this, the load on the storage system will be of a very different nature in an XProtect VMS compared to a standard file shares in an IT system. If this is not understood and properly accounted for during the design of the storage system for the XProtect VMS, it will result in performance issues and loss of recorded media data.

Standard IT storage system usage

Looking at the storage part of a standard distributed IT system, it typically consists of several servers and/or NAS's hosting file shares from an underlying storage system as well as users and other servers/services accessing (reading/writing/deleting) the files stored on the file shares.

When files in a standard IT system are written from e.g. a user's PC to a file share, it is not especially critical how fast the underlying storage system is from a data viewpoint, or in other words, how much time it takes to write the data. The reason for this is that the data that should be written is buffered at the source (for instance a user's PC) while being written to the storage system.

This means that data is not lost even if it takes some time writing it. Furthermore, the data is only transferred between source and destination when the destination is ready

to receive more data. So even though a slow storage system can be annoying for users, data is never lost.

Furthermore, only a smaller number of files are typically being read or written on the storage system at the same time, making the read/write operation mostly sequential, which gives a better performance for most storage technologies.

XProtect VMS storage system usage

Writing XProtect VMS media data to a storage system is very different compared to writing files in a standard IT system. This is mainly because in the XProtect VMS, cameras and other devices constantly produce new video, audio and metadata in real time, which needs to be sent to the XProtect VMS where it is recorded and viewed live.

Should the recording server and/or storage system be too slow to handle the media data in real-time, the standard IT transmission control and buffering mechanisms cannot be used to buffer, pause or throttle the transfer of data from the cameras, as this would create gaps in the recordings and cause video and audio to be paused and delayed when viewed live in the XProtect VMS clients.

In addition to the media data being received in real-time and not being able to be buffered or paused, the recording server typically records media data from many sources at the same time. This means that the media data from every cameras, microphones, speakers and metadata sources on the recording server is written in parallel to multiple files on the storage system.

Writing multiple files in parallel is called non-sequential or random writing and causes most standard storage systems to perform much slower compared to the sequential writing which is commonly used in standard IT systems, where only one or a few files are written at a time.

Finally, in extension to the process of recording the media data, the media data is also deleted again after the set retention time, or alternatively archived to another storage system, and then deleted after the set retention time.

Summary

The nature of live XProtect VMS data and IT data is very different, and thus the performance requirements are also very different which needs to be considered when choosing a storage system for the XProtect VMS.

If the XProtect VMS system designer does not take the real-time nature of the VMS data into account, but instead dimensions the storage system based on standard IT needs, it will result in issues in the XProtect VMS with missing recordings and slow playback performance.

Storage architecture

The Milestone XProtect VMS media database and storage architecture have been designed to provide the best possible performance, reliability, security and flexibility. This gives the XProtect VMS designer or administrator the freedom to choose a storage technology and architecture that fit the specific system needs, should the needs be performance, reliability, maintainability, size, cost or a combination of these.

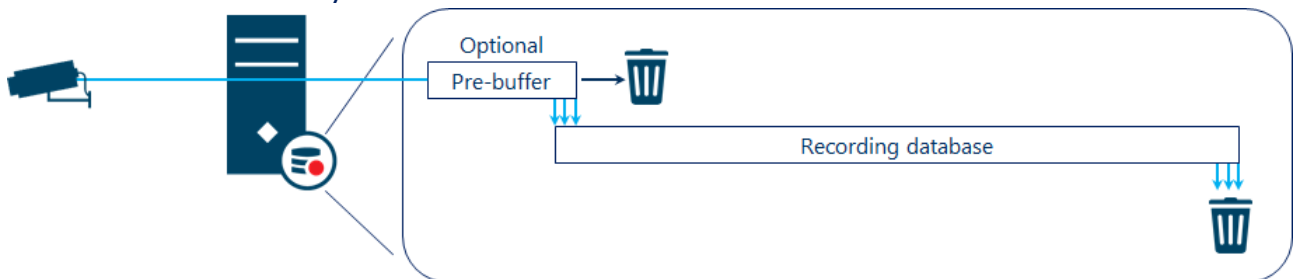
Storage and Media Databases

In the XProtect VMS the recording and archive parameters are configured as a part of an overall storage configuration, which specifies the retention time, size, encryption etc. for the recording database and any optional archive databases.

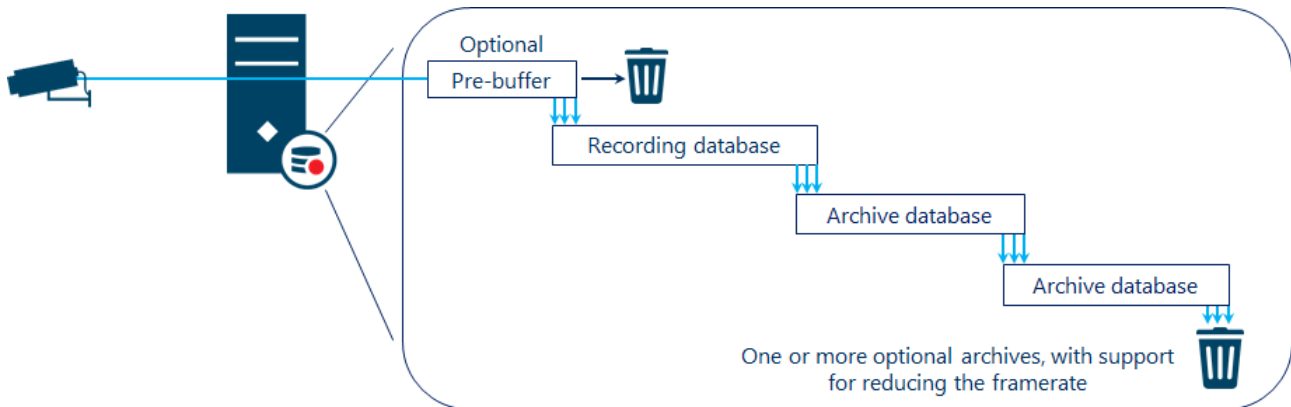
Each device assigned to use the defined storage will have its own database in it where the recorded media is stored. This means that there are as many media databases in the XProtect VMS as there are enabled devices.

The media is continuously streamed over the network from the device to the recording server, where it is initially stored in the (optional) pre-buffer, which when enabled, provides the ability to record some time before a recording triggering event occurs. See the 'Pre-buffer' section for details.

Once recording is triggered, the recordings are moved from the pre-buffer and written in the recording database. If no recording is triggered the data older than the set pre-buffer time is deleted. Once the data stored in the recording database is as old as the set retention time they are deleted.



In extension to the recording database, the storage definition supports a function called 'archiving', which is the process of moving the recordings with a set interval from the recording database to an archive database which is often stored at a different location. The archiving process can be repeated if needed. The archiving process also supports reducing the framerate of the recordings. See the 'Reduce Framerate' section for details.



No matter if pre-buffering and/or archiving are used, playback of recordings are completely transparent for the users of the XProtect VMS. The desired recordings are simply read by the media database from the location they are currently stored in, should it be the recording database or an archive database.

Pre-buffer


The purpose of the pre-buffer is to allow the XProtect VMS to record audio, video and metadata for a period leading up to an event triggering actual recording.

When enabled, the pre-buffer can work in two modes:

- Memory based
- Disk based

With memory-based pre-buffering, the media data is stored in the server's RAM, and only written to the recording database on the storage system when recording is triggered. This reduces the load, wear and tear on the storage system.

The reduction is inversely proportional to the percentage of media being recorded. If the XProtect VMS for example is configured to record always there isn't any reduction as everything needs to be recorded. Whereas with VMD triggered recording and VMD 20% of the time, the load, wear and tear on the storage systems is reduced by 80%.



Memory pre-buffering is the default and preferred option but is limited to maximum 15 seconds. Should a pre-buffer longer than 15 seconds be needed, it must be disk based.

With disk-based pre-buffering all media data is written directly to the storage system where it is stored for the set pre-buffer time. Once the set pre-buffer time has been exceeded, the recordings are either kept or deleted depending if any recordings have been triggered for the period of which they cover. Because all media is written to the storage system, and potentially deleted again (in case no recording has been

triggered), the load on the storage system is permanently high when using disk-based pre-buffering.

Recording database

The recording database is where recordings are initially written and stored. The term recording is used because the database files on the disk are open and actively being edited when the live feed from the devices is recorded.

The recorded media stays in the recording database until the retention time is reached and is then deleted – unless an archive has been defined.

If an archive has been defined, the recordings are not deleted when the retention time is reached, but instead remain in the recording database until the archive process is started.

Because of this, the recording database will at times hold recordings older than the set retention time when archiving is enabled.

As an example, the recording database is configured to 7 days retention, and the archive is scheduled to run once a day. This means that when the archive process starts, the recording database will hold 8 days of recordings: 7 days that should remain in the recording database, and 1 day which is about to be archived.

As mentioned earlier, recording many devices in parallel and in real-time to the recording databases causes a lot of non-sequential writing on the disks and storage system. Designing a storage system that is fast enough to handle this and large enough to store the recordings for the required time can be expensive, which is where archiving can be used to reduce the cost without compromising on performance.

Archive database

Archiving is not a requirement, but it can be used to optimize the overall cost and performance of the storage system. This is done by allowing a multi-tier storage architecture where live media is initially recorded to the first tier, and then later moved to the second tier.

The first tier can use smaller, faster but per Gigabyte more expensive disks, optimized for the non-sequential recording of live streamed data. The second tier can use larger, slower and per Gigabyte cheaper disks, optimized for storing the recorded media for the retention time needed.

The reason the disks used for the archive database can be slower is that the archive process that moves the recordings from the recording database to the archive database does this with only a few cameras at a time. This makes the disk access mostly sequential, in which case even cheaper, larger and slower disks perform well.

As with the recording database, the recordings remain in the archive until the retention time is reached, where the recordings are then deleted – unless a second archive has been configured, in which case the recordings remain in the archive database until moved to the next archive.

Note: The functions to archive the recordings more than one time are only available in XProtect Expert and XProtect Corporate.

Storage technologies

Knowing that live media data is of a real-time nature, and that data is written to many files in parallel when recorded, it is important to choose disks and storage system technologies that best support this kind of usage.

This section will cover how standard storage technologies work, the advantages and disadvantages they have, as well as how they can be used in an XProtect VMS environment in the best way.

This section will not cover hybrid storage solutions, which often consist of a mix of flash memory and/or SSD disks combined with traditional hard disks and/or tape robots. There are simply too many variants with specific benefits and/or pitfalls to be generically covered in this whitepaper.

With that said, hybrid storage solutions may work and provide certain benefits for some XProtect VMS installations. Should such a storage system be considered for a project, it is recommended to contact Milestone Systems Sales for assistance regarding the specific hybrid storage solution.

This section will also not cover object and cloud storage as it is not natively supported in the XProtect VMS. The reason these storage technologies are not natively supported is because the benefits they normally offer cannot be utilized by the XProtect VMS.

Object storage is well suited for large amounts of unstructured data where the metadata tags attached to the data are used to identify and find the data again. This make the storage system and use of it simple, as data can be stored in a flat structure and found again by searching for the data's identifier or metadata tags.

However, in an XProtect VMS, data is highly structured with the media database keeping a detailed index over the data and controlling where data is stored, making a hierarchical file-based storage system more ideal. Furthermore, because data may not be accessed directly on the storage system, but instead be accessed through API's on the recording server – to ensure only users with the right permissions can view the data – the additional benefits that object storage may offer cannot be used, making standard storage solutions better suited for XProtect VMS use.

Cloud storage is not supported either for several reasons. First, data is recorded at a constant high bitrate, hence, it would require a lot of upstream bandwidth to store XProtect VMS data in the cloud. Second, the traditional benefit of providing direct access to the data stored in the cloud from anywhere in the world does not apply to XProtect VMS data, as all data can only be accessed through API's on the recording server – again, to ensure only users with the right permissions can view the data. Therefore, because the XProtect VMS permanently sends data to the storage at a high

bitrate, and because users cannot access the data directly from the cloud but would have to stream it back through the recording server instead, standard storage solutions are better suited when using the XProtect VMS.


Standard disk technologies and key characteristics

There are primarily two types of disk drives – spinning plates hard disk drives (HDD) and solid-state drives (SSD).

Hard Disk Drive

HDDs very much resembles an old-style turntable with a pickup reading data (music) from a rotating plate. Instead of using a pickup and a vinyl plate, an HDD has a stack of rotating magnetic plates and an actuator arm with multiple electronic heads for reading and writing data as magnetic information on the plates.

As with the turntable analogy, when data should be read or written, the arm must move to the right track and wait for the plate to rotate to the right spot in the track, where data is then either read or written.



While the head moves and the disk rotates to the right place, data cannot be either read or written. This time is known as 'access time'. The access time varies depending on the distance from the current location of the head and plates to the new location that should be accessed. 'Access time' is thus specified for HDDs as average access time.

HDDs come with various rotation speeds, typically ranging from 4,200 to 15,000 RPM and have variations in how fast they can move the head. Because of this the average access time varies a lot across different HDDs – typically the average access time is between 3-20ms.

Furthermore, HDDs vary in performance depending on where data is read or written on the plates. Performance is highest when data is stored close to the plate's outer edge, and lowest when data is stored close to the plates center. The reason for this is that tracks are longer closer to the edge of the plates and thus can hold more data per rotation.

HDD benefits:

- They can be very large. Currently the largest are up to 14 Tera bytes
- They have a relatively low per Gigabyte cost
- They have a good read and write performance when the data is accessed sequentially, like reading or writing a single file at a time

HDD disadvantages:

- Due to having moving parts, they have a relatively lower durability when used extensively

- They have a relatively long average access time which has a huge impact when writing is non-sequential as it is when recording streamed data from multiple devices to multiple databases
- Disk performance varies depending on where on the plates the data is stored

So, in short, HDDs are slower (with non-sequential access), but larger and cheaper. This makes them best suited for storing the archive databases in an XProtect VMS.

With that said, HDDs can of course be used for recording the live data. However, much attention should be paid when choosing the disks to ensure they match the needs of the specific installation.

The test results below show the performance of an HDD. Notice that the write performance drops to ~8% of the maximum write performance when using non-sequential disk access.

Tests done:

1. Sequential – 64 Queues of data & 1 Thread
2. Non-sequential - 4 Queues of data & 4 Threads
3. Non-sequential - 16 Queues of data & 16 Threads
4. Non-sequential - 32 Queues of data & 32 Threads

	Read [MB/s]	Write [MB/s]
Seq Q64T1	76.47	63.91
4K Q4T4	1.009	5.566
4K Q16T16	0.929	5.055
4K Q32T32	0.915	4.795

Disclaimer: The test results do not illustrate absolute performance numbers of all HDDs. The results are included to illustrate the relatively large difference in performance between sequential (one file) and non-sequential (multiple files) disk access, and furthermore included to compare the performance with SSDs, which is covered below.

Solid-State Drive

SSDs have no moving parts, as they are completely made of flash memory chips mounted to a PCB.

Because they do not have any moving parts, the access time is much shorter - typically in the 0.01–0.1ms range.



SSD disks primarily come in two grades; Consumer-grade and Enterprise-grade;

- Consumer-grade SSDs are well suited for standard PC usage. However, due to having a lower write limit, it is not recommended to use these drives with an XProtect VMS as the permanent writing and deleting of data will wear out the drive up to ~25 times faster than Enterprise-grade SSDs
- Enterprise-grade SSDs are more expensive and can cost 2~10 times more compared to Consumer-grade SSDs. Enterprise-grade SSDs are thus not really relevant for standard PC usage. However, because of the various wear-leveling-techniques and increased write limit, these drives have a much better durability and are thus better suited for usage in an XProtect VMS

Enterprise-grade SSD disk benefits:

- They have a very short average access time
- Due to not having moving parts and having wear-leveling-techniques, Enterprise-grade SSDs have very high durability

Enterprise-grade SSD disk disadvantages:

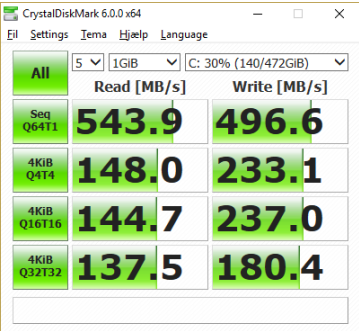
- They have a relatively higher per Gigabyte cost

So, in short SSD disks are very fast (with non-sequential access), but smaller and more expensive. This makes them best suited for storing the recording databases in an XProtect VMS.

The test results below show the performance of an SSD disk. Notice that the write performance is still ~36% of the maximum write performance when using non-sequential disk access

Tests done:

1. Sequential – 64 Queues of data & 1 Thread
2. Non-sequential - 4 Queues of data & 4 Threads
3. Non-sequential - 16 Queues of data & 16 Threads
4. Non-sequential - 32 Queues of data & 32 Threads



The screenshot shows the CrystalDiskMark 6.0.0 x64 interface. The test is configured for 5 tests, 1GB size, and C: 30% (140/472GB). The results are as follows:

	Read [MB/s]	Write [MB/s]
All	543.9	496.6
Seq Q64T1	543.9	496.6
4K Q4T4	148.0	233.1
4K Q16T16	144.7	237.0
4K Q32T32	137.5	180.4

Disclaimer: The test results do not illustrate absolute performance numbers of all SSD disks. The results are included to illustrate the relatively low difference in performance between sequential (one file) and non-sequential (multiple files) disk access, and furthermore to compare with HDDs.

One thing to remember though if wanting to use SSD disks, is to calculate the expected lifespan of the selected drives with the specific XProtect VMS usage. Typically, the disks' lifespan is described as the amount of data written per day to allow the drives to work for the entire warranty period. Calculating how much data is written per day allows an estimation of how long the SSDs can be expected to last.

RAID - Redundant Array of Independent Disks

In an XProtect VMS as in most IT systems, having just a single disk to store the data is typically not enough, as it doesn't provide the performance and size needed.

One option is to use JBOD and manually ensure that the load on the disks and data to be stored are distributed across the individual disks. However, while JBOD is a cheap solution that will work in any PC or server without extra controllers or software, it is not the most optimal solution as the storage size and performance across the disks cannot be combined and utilized in the most optimal way.

For more information about JBOD see the following Wikipedia page:

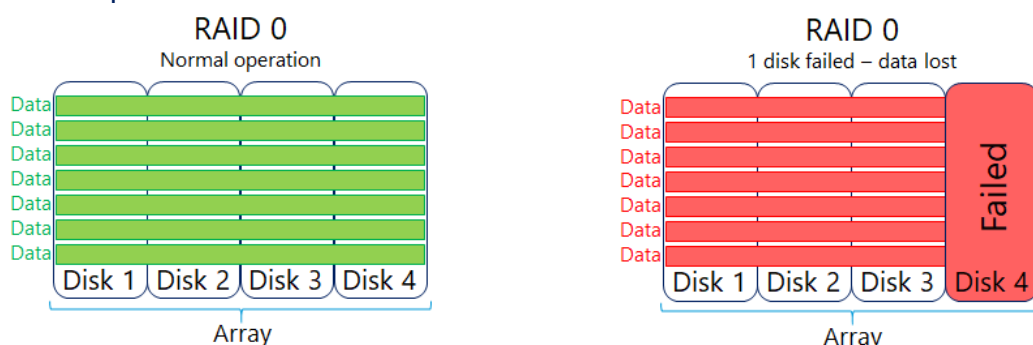
https://en.wikipedia.org/wiki/Non-RAID_drive_architectures#JBOD

An alternative to using JBOD is to use RAID technology, which can make the individual disks act as one large disk with higher performance and larger size. The most common RAID configurations are covered below.

RAID 0

The simplest form of RAID is called RAID 0 or 'striping'. With RAID 0 the disks (4 in below example) are combined in an array to form a single large drive. When data is written to this array, it is split into equally sized parts and each part is then written to the disks.

This gives excellent performance since the process of splitting the file in the RAID controller is simple, and since all drives write a part of the file in parallel. It effectively gives a write performance that is the sum of the individual disks.



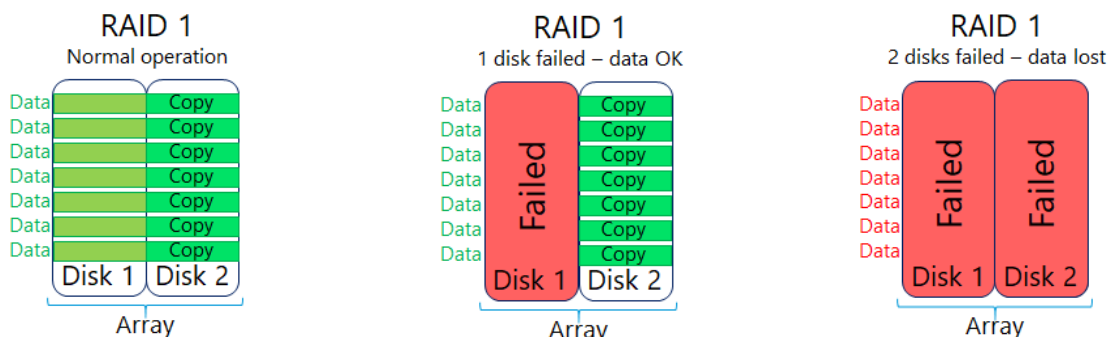
However, RAID 0 have one significant drawback – if one disk fails all data across all disks in the array is lost and cannot be recovered. This shortcoming can be addressed by using other RAID configurations.

For more information about RAID 0 see the following Wikipedia page:

https://en.wikipedia.org/wiki/Standard_RAID_levels#RAID_0

RAID 1

RAID 1, which is also known as 'mirroring', works by having a secondary drive act as a backup to the primary drive. Data is then written to both the primary and secondary drives effectively ensuring a 1:1 copy of the data.



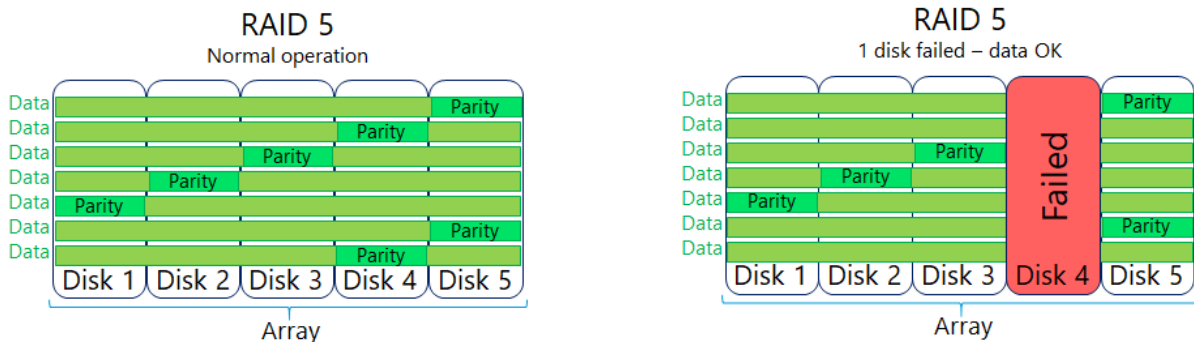
While RAID 1 protects the data and doesn't impact performance of the disk, it however requires twice the number of disks. Furthermore, it is not possible to make the array larger than the size of a single disk.

For more information about RAID 1 see the following Wikipedia page:

https://en.wikipedia.org/wiki/Standard_RAID_levels#RAID_1

RAID 5

RAID 5, which is also known as 'striping with parity', requires at least 3 disks. The disks are joined together in an array to form a single drive, much like with RAID 0. However, to protect against disk failures the data is stored across the disks in the array with some extra information called parity.



Should a disk fail this parity information can be used to reconstruct the missing data and thus allow the files to be accessed even if a disk has failed or is completely missing.

However, compared to RAID 0, the data protection comes at a cost of needing one more disk to achieve the same storage size. Furthermore, because parity needs to be calculated/recalculated for data being written/edited, RAID 5 is not as fast as RAID 0 using the same number of disks (plus 1 for parity).

With standard IT usage, a cache function in the RAID controller can alleviate a lot of the performance degradation. However, with XProtect VMS usage the media data is permanently received in real-time and in a non-sequential way appended to files already written to disk - triggering constant recalculations of the parity information.


This makes the cache function much less effective compared to standard IT data usage, and often something that is overlooked or not fully understood. If not taken into account, it can cause trouble in the VMS installation with missing recordings and slow playback experience in the VMS clients. Because of this, when considering using RAID 5 on the storage system for the VMS's recording database, it is important to take RAID 5's performance degradation with real-time non-sequential VMS data into account.

When considering RAID 5 on the storage system for the XProtect VMS's archive database, the performance of RAID 5 with VMS archive data compared to standard IT system data is more or less the same. This means that standard guidelines for storage performance known from IT usage can be used to calculate the needed performance when used for VMS archive databases.

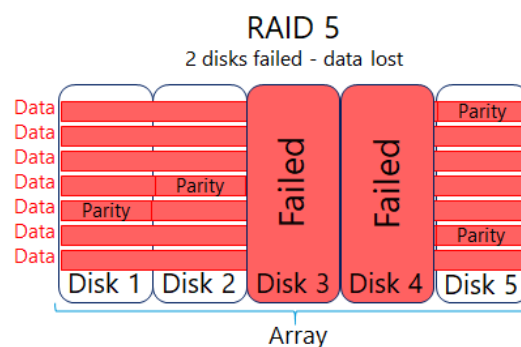
One more thing to consider with RAID 5 is the additional performance degradation of the array while a disk is in the failed state or being 'rebuild' - once the failed drive has been replaced.

When a disk is in the failed state, the RAID controller must work harder as data cannot be read directly from the disks since data from the failed disk is missing. Instead the data is reconstructed on-the-fly in the RAID controller by reading the data from the working disks and the extra parity information. This however, slows down performance and could impact XProtect VMS operation.

Once a failed drive is replaced, the RAID controller will rebuild the missing information on the replaced disk. This rebuild function puts an even greater load on the disks as all data from the remaining disks needs to be read and missing data recalculated and written to the new disk. This degradation in performance must be considered to ensure it does not impact the XProtect VMS performance.



Finally, just to cover this scenario as well, if two disks fail at the same time, all data on the array is lost and cannot be recovered. So, it is important that the disk array is monitored and that disks are replaced immediately if they fail.

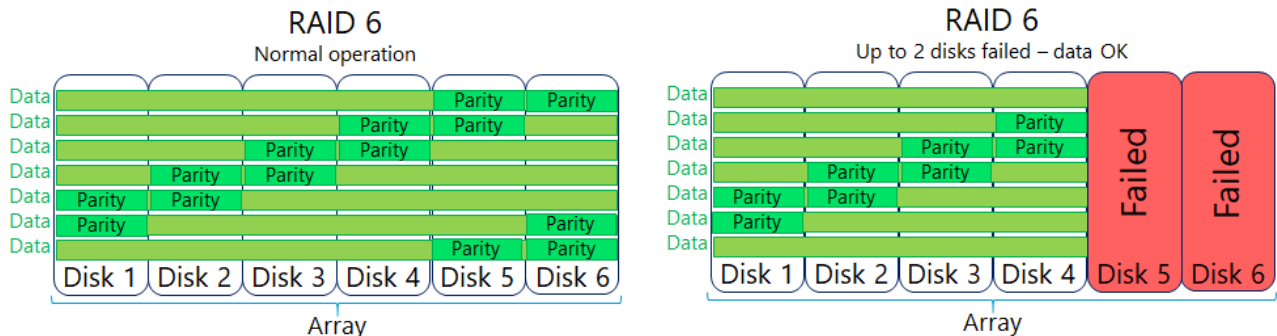


For more information about RAID 5 see the following Wikipedia page:

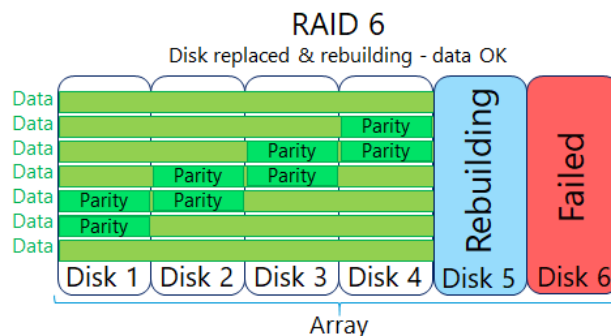
https://en.wikipedia.org/wiki/Standard_RAID_levels#RAID_5

RAID 6

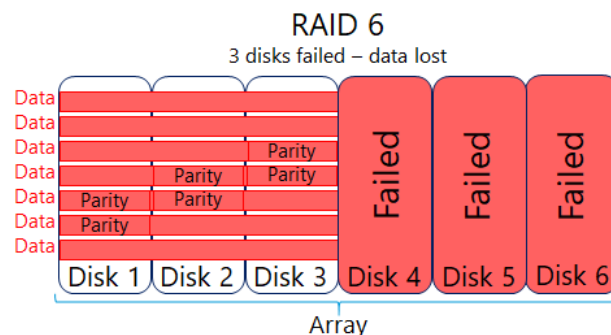
RAID 6 is basically an extension of RAID 5 where a second set of parity information is added. This allow the disk array to continue to work even if two disks have failed.



As with RAID 5 the performance of a RAID 6 array degrades if there are failed disks, and it degrades furthermore when replaced disks are being rebuild.



Finally, to also cover this scenario, if a third disk fails, data is lost and can't be recovered.



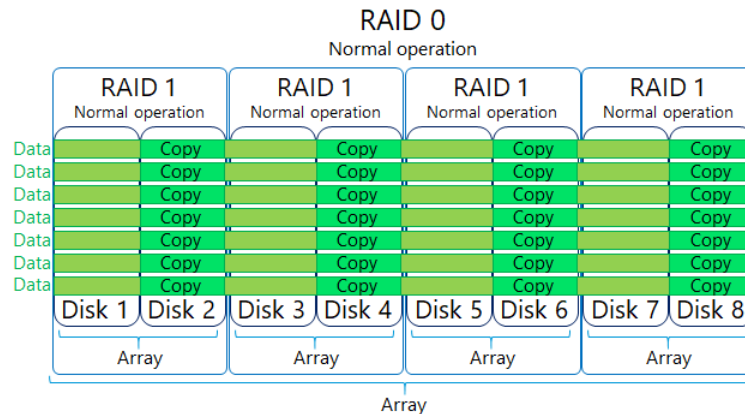
Recommendation for using RAID 6 with the XProtect VMS is the same as for RAID 5.

For more information about RAID 6 see the following Wikipedia page:

https://en.wikipedia.org/wiki/Standard_RAID_levels#RAID_6

RAID 10 (also called RAID 1+0 or RAID 1&0)

RAID 10 is a so-called nested RAID configuration as it utilizes two RAID levels at the same time. Sets of two (or more) disks are joined in several RAID 1 arrays. These RAID 1 arrays are then joined in a RAID 0 array.

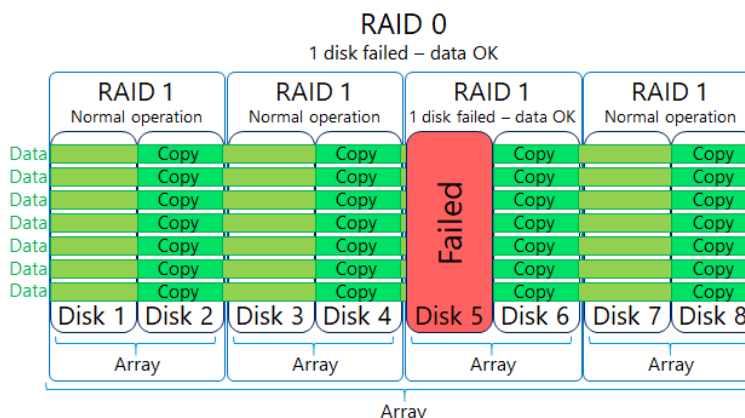


RAID 10 provides the best performance and redundancy and is often called 'the best RAID configuration for mission critical applications and databases'. In an XProtect VMS RAID 10 is especially good for the recording database – whether or not archiving is used.

However, the benefits of RAID 10 though come at the cost of needing a lot of hard disks making this configuration quite costly - especially when needing a lot of disk space as the XProtect VMS often requires.

This means that, for larger XProtect VMS systems with many cameras per recording server storing recordings for a longer period, RAID 10 may become too expensive. In this case a combination of a smaller RAID 10 array for the recording database, and a larger RAID 5 or 6 array for an archive database could be a more cost-efficient solution.

Should a disk fail, the data from the failed disk is simply read from the other disk in the RAID 1 pair. When a disk is failed the performance does not degrade.



The diagram illustrates the state of a RAID 1 array during a rebuild. It is divided into four sections, each representing a RAID 1 unit with two disks:

- RAID 1 Normal operation:** The first two disks (Disk 1 and Disk 2) are active and contain data. The label "Normal operation" is present.
- RAID 1 Normal operation:** The next two disks (Disk 3 and Disk 4) are active and contain data. The label "Normal operation" is present.
- RAID 1 Rebuilding - data OK:** The fifth disk (Disk 5) is in a "Rebuilding" state, indicated by a vertical blue bar with the word "Rebuilding" written vertically. The sixth disk (Disk 6) is active and contains data. The label "Rebuilding - data OK" is present.
- RAID 1 Normal operation:** The last two disks (Disk 7 and Disk 8) are active and contain data. The label "Normal operation" is present.

On the left, a vertical list of "Data" entries is shown, each corresponding to a row of data across the disks. The disks are labeled "Disk 1" through "Disk 8" at the bottom. Brackets below the disks group them into four pairs, each labeled "Array". A larger bracket at the very bottom spans all eight disks and is labeled "Array".

The diagram illustrates a RAID 1 configuration using 8 disks, organized into four pairs. Each pair is labeled 'RAID 1' at the top. The status of each RAID pair is indicated below the label: '1 disk failed – data OK' for the first and third pairs, and 'Normal operation' for the second and fourth pairs. A large red vertical rectangle labeled 'Failed' is positioned over Disk 2 in the first RAID pair and Disk 5 in the third RAID pair. To the left of the disks, the word 'Data' is written vertically, with horizontal lines connecting it to each of the eight data rows. Each data row contains eight green rectangles, one for each disk. In the first and third RAID pairs, the green rectangle for the failed disk (Disk 2 and Disk 5) is red. Below the disks, the labels 'Disk 1' through 'Disk 8' are placed under their respective columns. Brackets at the bottom group the disks into four 'Array' units, each containing two disks. A larger bracket at the very bottom groups all four arrays under the label 'Array'.

RAID 1
2 disks failed – data OK

RAID 1
1 disk failed – data OK

RAID 1
Normal operation

RAID 1
1 disk failed – data OK

RAID 1
Normal operation

Data

Failed

Failed

Disk 1 Disk 2 Disk 3 Disk 4 Disk 5 Disk 6 Disk 7 Disk 8

Array Array Array Array

Array

RAID 1
2 disks failed - data lost

RAID 1
Normal operation

RAID 1
Normal operation

RAID 1
Normal operation

Data
Data
Data
Data
Data
Data
Data

Failed
Failed

Disk 1
Disk 2
Disk 3
Disk 4
Disk 5
Disk 6
Disk 7
Disk 8

Copy
Copy
Copy
Copy
Copy
Copy
Copy
Copy

Array
Array
Array
Array

Array

[https://en.wikipedia.org/wiki/Nested_RAID_levels#RAID_10_\(RAID_1+0\)](https://en.wikipedia.org/wiki/Nested_RAID_levels#RAID_10_(RAID_1+0))

Other RAID configurations

Apart from the most used RAID configurations covered here, there are some additional RAID configurations called RAID 50, RAID 60 and RAID 100, which can offer some specific benefits.

For more information on these RAID configurations see the following Wikipedia page:

https://en.wikipedia.org/wiki/Nested_RAID_levels

RAID – hot spare

Hot spare is one or more extra disks attached to a RAID array in a storage system. With one or more hot spares available in the storage system, it can be configured to automatically replace a broken disk in the RAID array with one of the working hot spare disks.

This ensures that a disk is replaced immediately without manual intervention should a disk fail in the RAID array, allowing the RAID array to be rebuilt as fast as possible and minimize the time the RAID array is impacted by the failed disk.

Even though the storage system automatically replaces broken disks with hot spare disks, it is still important to monitor the storage system and replace the broken disks in a timely manner to ensure working hot spares are available continuously.

For optimal performance and availability of the storage system used in an XProtect VMS, it is recommended that the RAID array is configured with one or more hot spare disks.

RAID controllers

In addition to choosing disk type/model, number of disks and RAID level to obtain the performance and disk space needed, it is also very important to choose a RAID controller that can utilize the full potential of the selected RAID level and disks to deliver maximum performance.

Often the RAID controller is actually the bottleneck in the storage solution – especially if using software RAID defined in Windows, or if using cheap “no-brand” RAID controllers.

Because of this, it is very important to check the performance specification of the RAID controller to ensure that it can deliver what is needed at the desired RAID level, including during periods where a disk is in the failed state or the RAID array is being rebuilt after a disk has been replaced.

Storage enclosure technologies

The disks used for the RAID can of course be installed in the PC or server running the XProtect VMS recording server if it has space for it, but for larger installations needing a lot of storage and thus many disks, the disks are typically installed in an external storage enclosure. Such an external storage enclosure is typically either a DAS, NAS or SAN.

DAS

A DAS is an external enclosure with its own power supply and disks. A DAS can only be connected to a single server and the controller for the DAS is typically installed inside the server. Furthermore, the RAID array configuration is defined and stored in the server/controller.

A DAS is well suited for both recording and archive database usage with the XProtect VMS, and because it provides the performance and storage a single recording server needs, and typically is less expensive, it is most often the ideal choice for recording server storage.

For more information on DAS see the following Wikipedia page:

https://en.wikipedia.org/wiki/Direct-attached_storage

NAS

Like with a DAS, a NAS is an external enclosure with its own power supply and disks. However, unlike a DAS, a NAS is not directly connected or attached to the servers. Instead a NAS is connected to the standard IT network and provides access via one or more file shares, which servers and users can access simultaneously - if they have access permissions for the file shares.

In the past, a NAS was primarily a storage solution for home users and smaller businesses that needed a small, cheap and simple file share to store their files. However, over the last years, large professional NASs with high performance, high capacity and redundancy functionality have become available.

With the XProtect VMS a NAS may only be used for the archive databases. The reason for this is that the recording database needs uninterrupted block-level access to the storage system to ensure high performance and continuous recording of live media data. Because a NAS is connected to a standard IT network and uses file-level access, direct uninterrupted disk access cannot be guaranteed as needed, which means that even small delays or gaps in the communication will cause performance issues and loss of recordings.

When a NAS is used for the archive database, the data has already been recorded to the recording database. This means that when moving the recordings to the archive database having uninterrupted block-level disk access is no longer needed. File-level access is enough as data is not lost in case of an interruption. The archive job simply pauses and resumes once the connection has been restored.

One thing to remember when using a NAS for the archive database, is to secure access to the NAS's file share that stores the media files so only the XProtect VMS recording server can access the file share. This is needed because no one else other than the XProtect VMS recording server should be able to access the media database files. All other systems or users that should access the media data stored in the database files, must access it through the recording server, which then ensures that only systems or users with the correct permissions get access to the media data.

For more information on NAS see the following Wikipedia pages:

https://en.wikipedia.org/wiki/Network-attached_storage

SAN

A SAN, as the two previous types, is an external storage enclosure with its own power supply, storage controller and room for multiple disks. The SAN provides access to the defined disk arrays to one or more servers. A single server is typically connected to the SAN via a direct optical connection called a Fibre Channel. If more servers share the SAN, the optical cables are connected to a Fibre Channel switch, which then is connected to the SAN. Alternatively, iSCSI can be used to provide the connection to the SAN over the standard IP network. If iSCSI is used a dedicated network is typically recommended for performance reasons.

Each defined array in the SAN is assigned to a specific server. This means that even though two servers may share the SAN, the data on the individual disk arrays in the SAN is not shared between the servers but belongs to a single specific server and can only be accessed through that server.

A SAN is typically a more advanced and expensive storage solution compared to DAS and NAS, and it is mostly used in installations where more servers share the storage infrastructure and the flexibility that the SAN offers.

A SAN is well suited for both recording and archive database usage with the XProtect VMS. However, because a single XProtect VMS recording server can/will utilize the entire performance and storage space on a SAN, some of the benefits like sharing the SAN across multiple servers cannot be utilized just as well in a XProtect VMS installation as in standard IT installations, making it a more expensive solution.

For more information on SAN see the following Wikipedia pages:

https://en.wikipedia.org/wiki/Storage_area_network

Storage configuration

This section provides an overview of how to configure storage with optional archives.

For detailed information on how to configure all storage and recording parameters, please refer to the 'XProtect VMS - Administrator manual' which can be found here:

<https://www.milestonesys.com/support/help-yourself/manuals-and-guides/?prod=3&type=13&lang=27>

Basic configuration

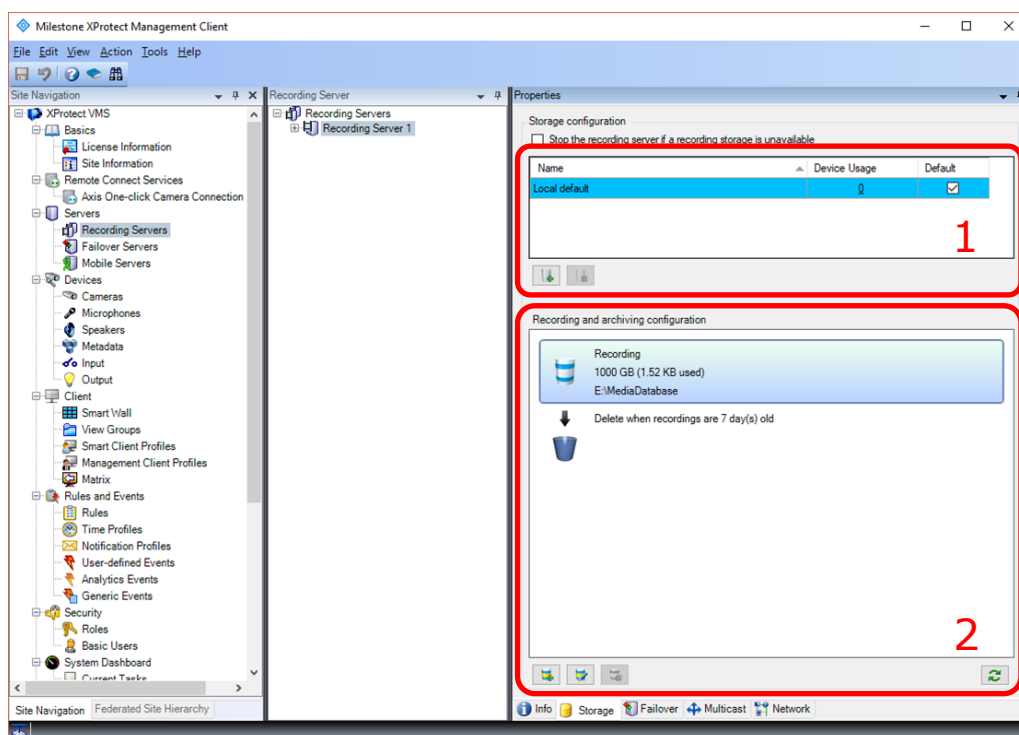
Configuration of the storage path, retention time and optional archiving is done per recording server by using the XProtect Management Client. Once the Management Client has been logged in to the XProtect VMS, select the *Recording Server* node in the navigation tree. Then select the recording server to manage in the recording server tree, and finally select the *Storage* tab.

XProtect Storage Architecture and Recommendations

Once the storage tab has been selected, the default storage configuration is displayed. This default storage configuration (path and retention time) was set in the installer when the recording server was initially installed.

The storage configuration consist of two areas.


- Area 1: Lists the defined storage configurations for the selected recording server. Furthermore, it shows how many devices currently use each storage configuration and which storage configuration is the default for new devices added to the recording server
- Area 2: Shows the storage path, retention time and currently used space for the recording database and any defined archives for the storage configuration selected in Area 1




To edit the settings for the selected storage configuration, click the storage graphics or the *Edit Recording Storage* button  in area 2.

One thing to notice is the *Retention time* and *Maximum size* settings.

These settings specify how old recordings need to be or how much space they may use (whichever comes first) before they are deleted or alternatively moved to an archive at the next archive time.



Archiving

If recordings should be archived to another location, click the *Add Archive* button  to add an archive.

The retention and maximum size settings work in the same way as in the recording settings dialog where they control when recordings are deleted or moved to yet another archive.

One option to pay attention to when configuring an archive is the *Schedule* definition, which controls when recordings are moved from the previous recording location and into the archive.

Although the archive schedule can be set freely, from once an hour up to once a month, it is important to consider that recordings will remain in the previous storage location until the archive schedule triggers the recordings to be moved to the archive.

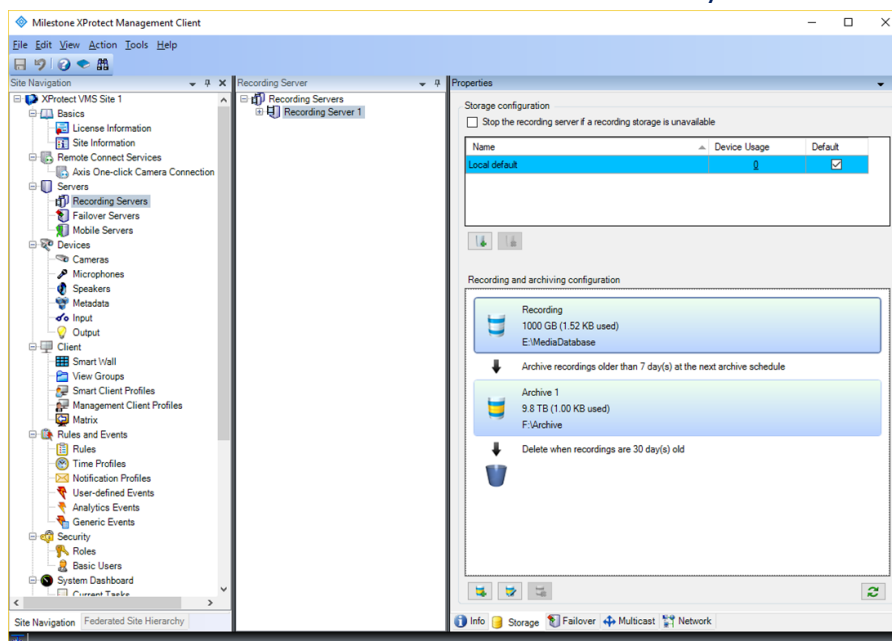
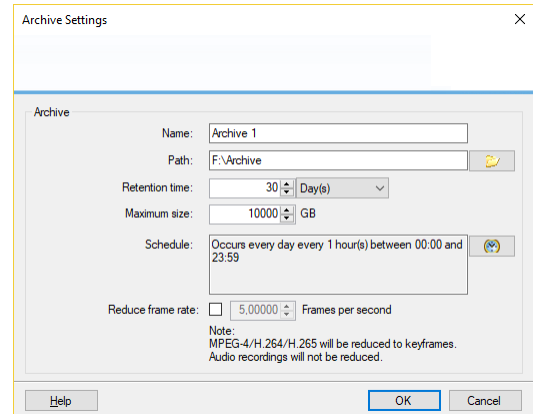
Example; if the retention setting on the initial recording location is set to 7 days and the schedule in the archive is set to archive once every day, there will be 8 days of recordings in the initial recording location when the archive starts.

To ensure the best performance, it is recommended to archive every 4 hours.

The same applies if subsequent archives are defined. Recordings will stay in the previous storage location until next archive is scheduled.

Note: The functions to archive the recordings more than one time are only available in XProtect Expert and XProtect Corporate.

Below screenshot shows a single storage configuration containing the initial recording location and one archive with a total retention time of 30 days.



Reduce framerate

If an XProtect VMS installation is required to record video at a high framerate, and store it for a long time, a storage system with a lot of space is needed - which can be expensive to purchase and maintain. However, as the recorded video gets older, it typically becomes less important and thus a lower recording framerate can sometimes be acceptable.

Reducing the framerate of older recordings when they are archived can reduce the space needed to store the video recordings. However, the amount of reduction that can be achieved depends on the video format used. Furthermore, the process of archiving itself creates an extra load on the recording server and storage system, which must be considered.

MJPEG video

When recording video in MJPEG format - which is basically a series of independent JPEG images - the saving in space is proportionate to the number of images the recordings are reduced by. Furthermore, because the MJPEG stream is just a series of JPEG images that can be decoded independently from one another, the framerate can be reduced to any framerate lower than the original framerate the video was recorded in.

Example: Recording video as MJPEG with 30 FPS and each image using 100KB.

- One second of the initial 30 FPS recordings will use $100\text{KB} * 30 = 3,000\text{KB}$
- When reducing the framerate to 5 FPS, one second of video will now use $100\text{KB} * 5 = 500\text{KB}$ - a saving in needed disk space of 83.33%

This means that when recording video in the MJPEG format there is a lot of space that can be saved by reducing the framerate of the recordings over time.

MPEG video

When video is recorded in the MPEG formats, the benefits of reducing the framerate of the recordings is not as big as with MJPEG. The reason for this lies in how these codecs compress the video.

In MPEG formats, the images are grouped into units called GOP's. Each GOP consists of an initial image called a 'keyframe' or 'I-frame' and a series of "partial frames", called 'B-frames' and/or 'P-frames', which only contain information about changes in the image compared to the previous/next images in the GOP.

The keyframe is similar to a JPEG and can be decoded independently. The subsequent images in the GOP only contain information about changes compared to the previous/next image in the GOP. The 'B-frames' and 'P-frames' can thus not stand alone and be decoded independently. They can only be decoded as part of a complete sequence starting with the keyframe.

Furthermore, because the keyframe is the reference image for decoding the rest of the images in the GOP, and the subsequent images only consist of updates to the

keyframe or previous images, the keyframe takes up a lot of the space in the GOP – especially if there is not much movement in the video. If there isn't much movement in the video, which is often the case in an XProtect VMS, the keyframe can use 60-80% of the entire GOP.

Additionally, because the 'B frames' and 'P frames' cannot be decoded independently, MPEG video can only be reduced to the keyframe interval (default 1 FPS) or less (e.g. every second keyframe). Moreover, because the keyframe uses most of the data in the GOP, reducing the FPS from 30 to 1 may result in saving only 20-40% of the storage used.

It is possible to configure a shorter GOP length, of 0.5 seconds for example, making a reduced framerate of 2 FPS possible. However, doing so is not recommended, as having two GOP's per second - and thus also two large keyframes per second - will increase the network and storage load and increase the disk space needed until the recordings are archived and reduced in framerate.

For more information about MPEG-4, H.264 & H.265 video encoding and GOP's see:

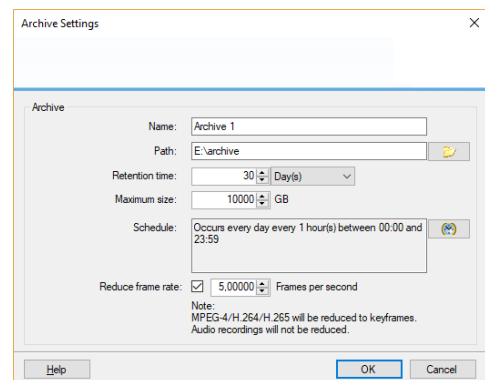
https://en.wikipedia.org/wiki/H.264/MPEG-4_AVC

https://en.wikipedia.org/wiki/High_Efficiency_Video_Coding

https://en.wikipedia.org/wiki/Group_of_picture

Configuration

The function to reduce the framerate is configured as part of the archive configuration. When enabled, only the set subset of the video frames are read from the previous media database and transferred to the archive.



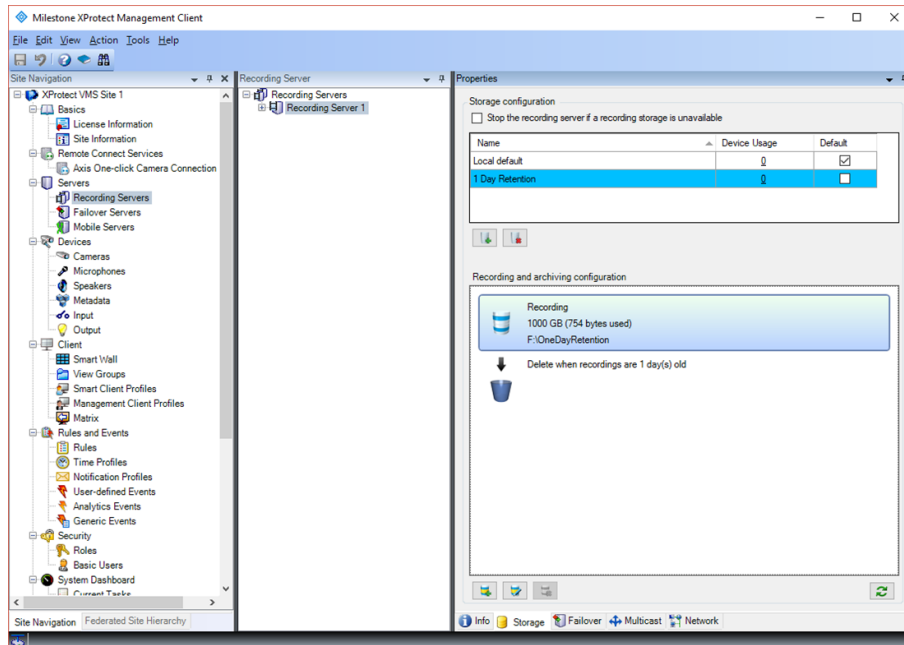
Note: The function to reduce the framerate during archiving is only available in XProtect Expert and XProtect Corporate.

Differentiated retention time

Should devices on the same recording server use different retention times, further storage configurations can be added by clicking the *Add Recording Storage* button



in area 1 of the dialog.



Above screenshot shows a recording server with two storage configurations defined ('Local default' and '1 Day Retention').

Once multiple storage configurations have been defined, the one to use for a specific device can be selected in the device's *Storage* tab.

Low disk space handling

The recording server supports a last resort fallback function to handle unexpected low disk space situations. If triggered, the recording server tries to make space for new recordings or archives by either moving or deleting exiting recordings depending on the configuration and the drive that runs low on space:

- If a next archive location relative to the drive that runs low is defined, the oldest recordings/archives are automatically archived to this archive location regardless of the defined archive schedule
- If no archives are defined, or if it is the drive storing the final archive location that runs low on space, the oldest recordings are deleted

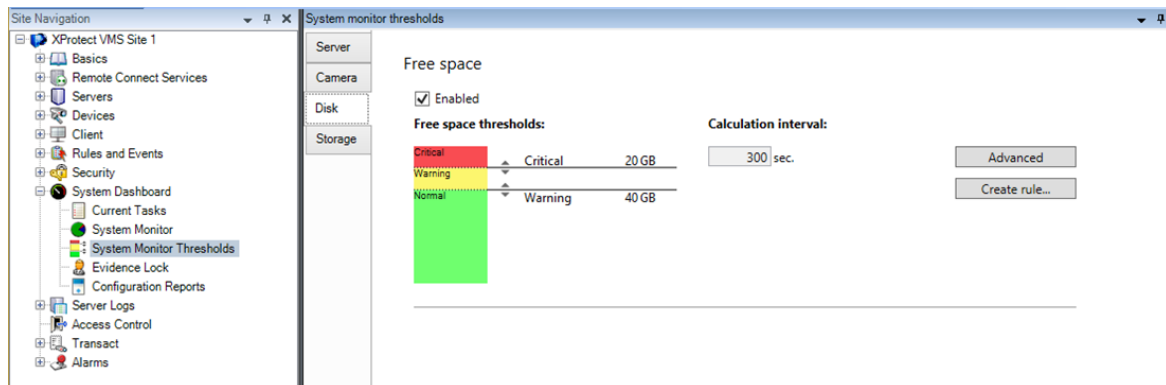
However, the success of the function to archive the recordings to the next location, depends on whether the drives can move the recordings from the drive with low disk space to the next archive drive faster than new recordings are being stored on the disk with low disk space. If this cannot be done fast enough, the disk will eventually become completely full which will trigger recordings to be deleted.

Because there is a small risk that recordings will end up being deleted, it is not recommended to design the recording server and storage around utilizing this last resort fallback function as part of the regular operation. This function should only be used as a last resort to handle a situation where the system is running out of disk space.

System monitor events

To catch low disk space situations before they actually happen and trigger above functionality, it is recommended to configure the low disk space thresholds, which can be found in the *System Monitor* to appropriate levels.

When defined, events will be triggered when the free disk space becomes lower than the set warning and critical levels. These events can then be used to trigger alarms or send email notifications via the rule system.



Securing XProtect VMS media databases

Permissions

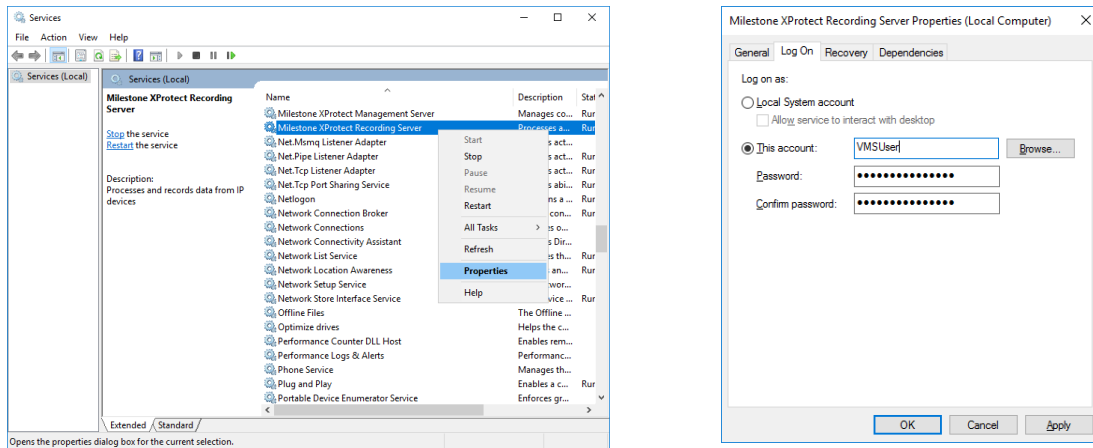
Access to live and recorded audio, video and metadata through the recording server's APIs, for instance when using the XProtect VMS product clients or MIP SDK, is protected by user authentication and security permissions set for the users in the roles.

To protect the media database files on the recording server disk or NAS/file share from unauthorized direct access, it is recommended to ensure that only IT administrators and/or XProtect VMS administrators have login permissions to the server running the recording server. Furthermore, if the recording server archives recordings to a NAS/file share, it is also recommended to configure security and permissions control for accessing the NAS/file share where the media database is stored. This ensures that only the recording server can access the files and not any user with access to the network.

When enabling security on a NAS/file share, the recording server service must be changed from running under the *Local System account*, to running under a Windows Domain or Workgroup account. The account the recording server service is set to run

under must then be granted rights to access the NAS/file share holding the media database.

The account to run the recording server service under is set via the Microsoft Windows 'Services' dialog.

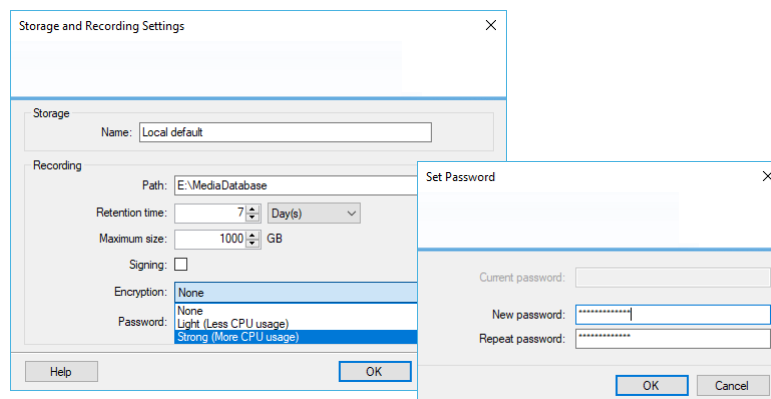


Open the *Services* dialog in Microsoft Windows, find the *Milestone XProtect Recording Server* service in the list. Right-click and select the *Properties* option. In the shown dialog, select *Log On*, then *This account* and enter user credentials for a user with access to the file share.

Media database encryption

In extension to protecting access to the server running the recording server and protection access to the NAS/file shares, it is also possible to add a layer of security on the media database itself by enabling the media database to encrypt the files stored on the disk.

Encryption is enabled for a storage definition by selecting *Light* or *Strong* in the *Encryption* dropdown in the 'Storage and Recording Settings' dialog and entering a password.



It is important to document and keep the password in a safe place as it may be needed later should the settings be changed, or the media databases need to be opened and read directly by the 'XProtect Smart Client – Player' for example to view the recordings in a backup or a copy of the media database.

The encryption standard used for both the *Light* and *Strong* options are AES-256. The difference between the two encryption options lie in the amount of data that is encrypted.

Light encryption

- *Light* encryption only encrypt a smaller section in the beginning of each record (a JPEG image, a video GOP, an audio segment or a metadata file) stored in the media database. Without access to this initial section of the record, it is impossible or at least extremely difficult to decode/understand the remaining data in the record
- Because only a small part of the record is encrypted it requires less processing power to encrypt the record

Strong encryption

- *Strong* encryption encrypt the entire record stored in the media database
- Because all data in the record is encrypted it requires more processing power to encrypt the record

Note: Media database encryption is only available in XProtect Expert and XProtect Corporate.

Media database signing

In addition to protecting the recordings with permission settings and encryption, it is also possible to add a digital signature to the media database. With a digital signature added to the media database it is possible to check and verify that the original media database files on the recording server disk or NAS/file share have not been tampered with.

When exporting in the XProtect Smart Client and using the XProtect format, it is possible to include the original digital signature, as well as to add a second signature during the export. When this export is viewed in the 'XProtect Smart Client – Player', the two digital signatures can be verified to check that the original recordings and exported data have not been tampered with.

Note: Media database signing of the media databases during recording is only available in XProtect Expert and XProtect Corporate. Adding a signature during an export is available in all XProtect VMS.


Evidence Lock

XProtect Corporate includes a function called 'Evidence Lock' that can protect important recordings from being deleted at the set retention time.

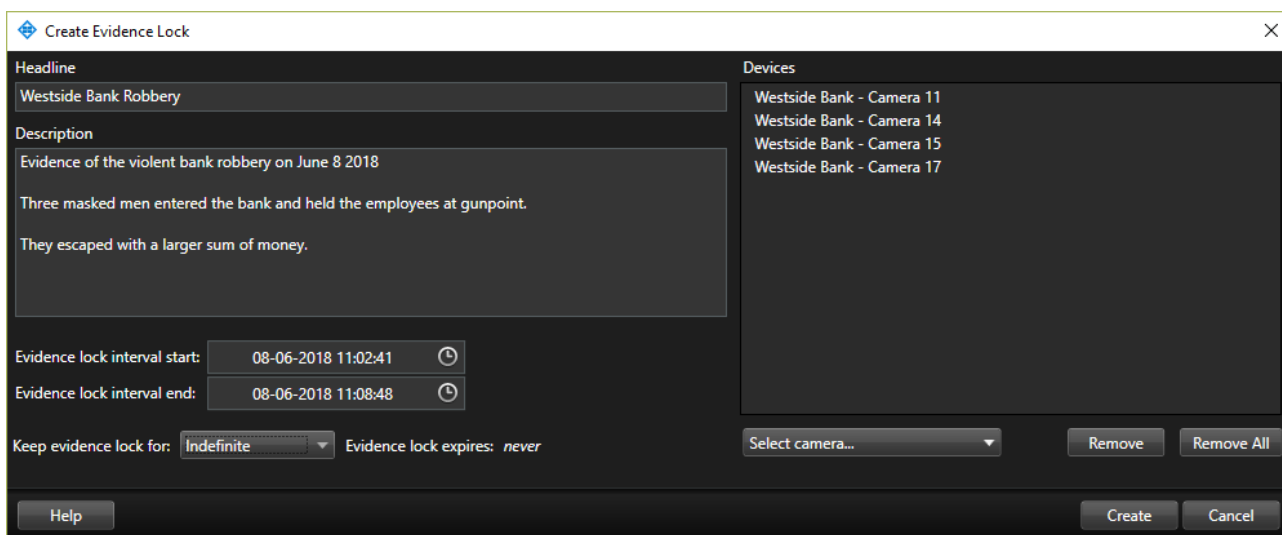
Using Evidence Lock

Evidence Locks are created using the XProtect Smart Client by; going to *Playback*, selecting a time period to protect via the timeline, checkmark the cameras in the view

to include, and finally clicking the *Evidence Lock* button and selecting *Create...*

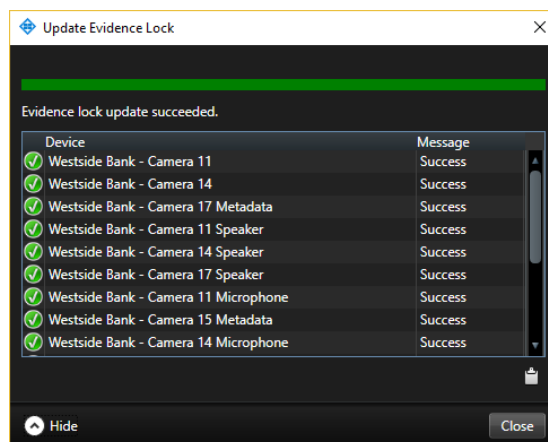


Doing so will present the below dialog, where the Evidence Lock details can be set - including how long to protect the recordings from being deleted by the normal retention time settings.




Using the XProtect Management Client, the *Keep evidence lock for:* dropdown options can be configured to only present the user with the Evidence Lock retention times approved for the specific installation and/or user

Once the Evidence Lock has been created, a status dialog is displayed showing which devices are included in the Evidence Lock – including related devices - and if the action has been successful.

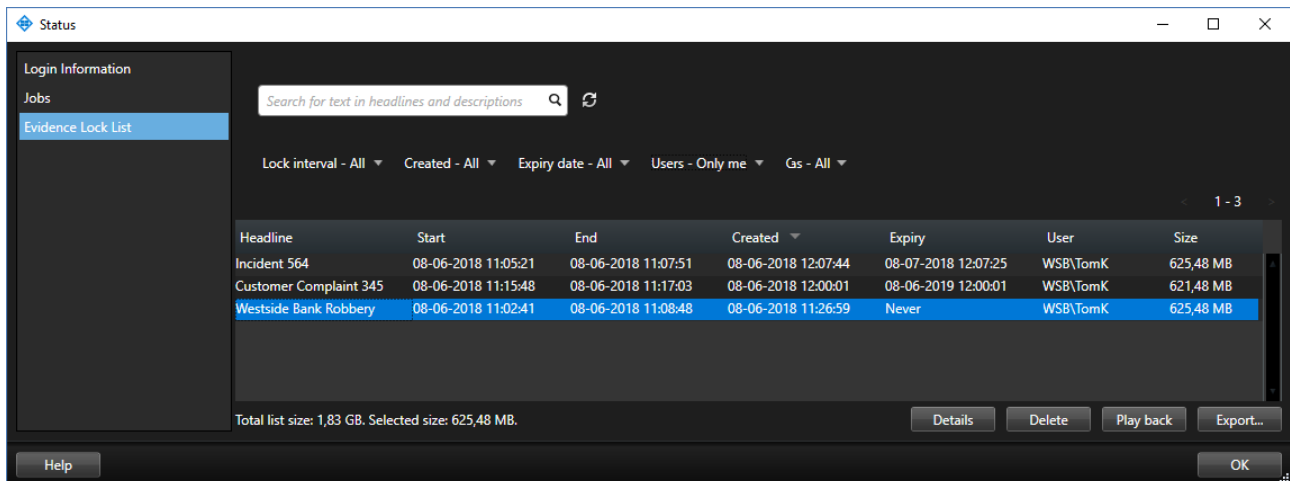


To see an overview of the Evidence Locks in the system using the XProtect Smart Client, do one of the following;

- Click the *status* button in the top right corner of the XProtect Smart Client, and select *Evidence Lock List*
- Click the *Evidence Lock* button on the *Timeline* in the *Playback* tab and select *View...*



Doing so will display the below dialog showing the Evidence Locks in the system created by the current user.



When viewing the list of Evidence Locks using the XProtect Smart Client, it is possible to filter and search for the Evidence Locks in the system, including Evidence Locks created by other users.

Once an Evidence Lock in the list is selected, it is possible to directly playback or export the recordings. Furthermore, if the user has the security permissions for it, it is possible to edit all parameters of the Evidence Lock and add or remove devices from the Evidence Lock.


Finally, if the user has the security permissions for it, the Evidence Lock can be deleted. Deleting an Evidence Lock does not necessarily mean that the recordings are also deleted as this depends on whether they are younger or older than the retention time set for the storage configuration on the recording server. If recordings are older than the set retention time, the recordings will also be deleted when deleting the Evidence Lock.

In addition to viewing Evidence Locks using the XProtect Smart Client, the XProtect Management Client can be used to provide a more detailed overview of the Evidence Locks in the system. The XProtect Management Client cannot however manage the Evidence Locks in the system, only the XProtect Smart Client can be used for that.

Storage impact


When creating an Evidence Lock, the recordings in the media database for the selected devices are marked as protected by an Evidence Lock. This is done to let the recording server and media database know how to handle and properly protect these recordings – even in a startup or offline scenario where the recording server might not be able to communicate with the management server.

In order to avoid tampering with possibly encrypted and signed media databases, as well as for performance reasons, an Evidence Lock will always lock and protect at least a single 1 hour segment of the media database for each device – no matter how small the time span of the actual Evidence Lock may be.



XProtect Storage Architecture and Recommendations

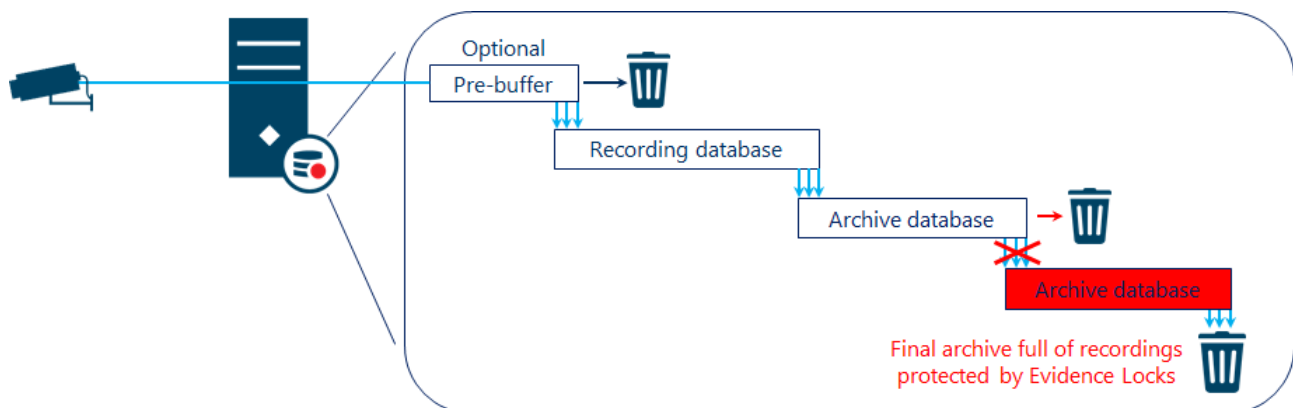
If the selected time span for the Evidence Lock spans multiple one-hour segments in the media database, all segments it spans are locked and protected.



This means that every single Evidence lock will protect more recordings than just the selected time span. For short sequences, like for instance an Evidence Lock covering a few minutes, this overhead may percentage-wise be quite large. For longer sequences, like for instance an Evidence Lock spanning 24 hours, the overhead will percentage-wise be much smaller.

If archiving is enabled, the recordings protected by the Evidence Lock will be moved to the next archive just as all other recordings, however they are archived without reducing the framerate – if reduce framerate has been enabled for the archive.

In case the final storage location in a storage definition becomes full with recordings protected by Evidence Locks, the recordings in the previous storage location are deleted instead of archived to make room for new recordings.



If all storage locations in the entire storage definition become full with recordings protected by Evidence Locks, the recording of new media will stop completely.

To get recording started again, some Evidence Locks must be deleted, or more disk space added.

Furthermore, to prevent this situation from happening it is recommended not to protect recordings indefinitely – or at least ensure this option is used with great care and focus on disk space usage. An alternative to protecting recorded media forever by using Evidence Lock, is to export them to external media for safe keeping.

Regardless of Evidence Lock retention settings, it is recommended to use the various disk space and archive events to trigger alarms and/or email notifications should the disk space becomes low.

Note: Evidence Lock is only available in XProtect Corporate.

Media database performance and recommendations

When designing or specifying an XProtect VMS system, there are many variables that impact the load and performance of the recording server and storage system, and thus the design of the XProtect VMS – for instance:

- Devices
 - Video resolution, video codec, framerate, bitrate, etc.
- Recording server
 - Recording server spec's
 - Number of devices per server
- Storage solution
 - Disks or combination of disks to use
 - RAID level
 - Use of archiving
- When and what to record
 - Record – always, record on motion, record on event, or a combination
 - Record – video, audio, metadata or a combination
- How many are expected to view recorded media at the same time

This section will provide some general guidelines to observe and follow when designing an XProtect VMS.

Recording database

When choosing a storage system for the recording database it would be optimal if the XProtect VMS requirements could be calculated, and the specific performance of the storage system could be looked up in a table or specification sheet – maybe as IOPS as this is often thought to be a good metric. Unfortunately, IOPS only makes sense if below is true:

- The size of the data block being written by the application matches the size used for the IOPS specified for the storage system
- The size of the data blocks being written is of a constant size
- The access time (or response time) is known – which depends on knowing the level of sequential/non-sequential disk access

In the real world, an XProtect VMS system and its usage of the storage system is much more dynamic which makes IOPS meaningless:

- Data blocks vary in size from large blocks of video data, to very small blocks of metadata
- In addition to the media data being recorded, the media database at frequent but variable intervals, makes small updates to index files
- Video, audio and metadata streams are not of a fixed size, they vary over time depending on what is being captured by the camera, microphone or sensor

- The disk access time varies over time depending on what type of data is being written at any given moment

This means that because the specific application load and storage performance cannot be known or calculated, an exact answer to what storage is needed cannot be given.

However, based on knowledge of how the XProtect VMS and media databases work, as well as results from tests and real-world experience, the below tables can be used to get a general idea of the recommended disk types as a function of cameras per recording server and bitrate/framerate.

Disk recommendations

The reason the 'per camera' units in the below tables are different for cameras using MPEG and MJPEG is that for MPEG devices, video is by default received as one GOP per second regardless of the framerate, and thus is stored as a single record in the media database per second no matter the framerate. So, with MPEG it is the bitrate that mostly matters and not the framerate.

With MJPEG each individual image is stored as a separate record in the media database. So, with MJPEG the framerate has a much higher impact on the disk performance than the bitrate.

Recording video only

Below tables show the general disk recommendations for various scenarios.

The following is assumed:

- Pre-buffer is running in memory
- Only video is recorded - no audio or metadata
- Video is only recorded when motion is detected on the video stream
- Recordings are not archived
- When using HDDs, they run at 10.000 RPM or more and have low average access time
- Storage is configured using RAID 1

Legend
Single HDD disk
Two HDD disks
Four or more HDD disks / 1-2 SSD disks
Several SSD Disks

MPEG-4/H.264/H.265 - 1-40% Motion						
Average bitrate per camera	Number of cameras					
	10	25	50	100	250	500
1 Mbit/s						
2 Mbit/s						
4 Mbit/s						
8 Mbit/s						

MJPEG - 1-40% Motion						
Framerate per camera	Number of cameras					
	10	25	50	100	250	500
1 FPS						
5 FPS						
10 FPS						
30 FPS						

MPEG-4/H.264/H.265 - 40-70% Motion						
Average bitrate per camera	Number of cameras					
	10	25	50	100	250	500
1 Mbit/s						
2 Mbit/s						
4 Mbit/s						
8 Mbit/s						

MJPEG - 40-70% Motion						
Framerate per camera	Number of cameras					
	10	25	50	100	250	500
1 FPS						
5 FPS						
10 FPS						
30 FPS						

MPEG-4/H.264/H.265 - 70-100% Motion						
Average bitrate per camera	Number of cameras					
	10	25	50	100	250	500
1 Mbit/s						
2 Mbit/s						
4 Mbit/s						
8 Mbit/s						

MJPEG - 70-100% Motion						
Framerate per camera	Number of cameras					
	10	25	50	100	250	500
1 FPS						
5 FPS						
10 FPS						
30 FPS						

Note: If enabling archiving, the recommendations in above tables can still be used by reducing the number of cameras listed in the tables to a half.

Average bitrate per camera	Number of cameras					
	10	25	50	100	250	500



Framerate per camera	Number of cameras					
	5	12	25	50	125	250

Recording video, audio and metadata

When enabling recording of audio and metadata as well, the load on the storage system increases – not so much because of the increase in total bitrate being recorded, but because each new audio and metadata device produces data that also needs to be stored on the disk. Because there are more devices, there will be an increase in non-sequential disk access which causes HDDs to perform slower.

So, when recording audio and metadata in addition to video, a faster storage system is needed – even though the bitrate percentagewise doesn't increase much.

Below tables show the general disk recommendations for various scenarios.

The following is assumed:

- Pre-buffer is running in memory
- A camera, a microphone, a speaker and a metadata source are combined and listed as a 'unit' in the tables
- Video, audio and metadata are only recorded when motion is detected on the video stream
- Recordings are not archived
- When using HDDs, they run at 10.000 RPM or more and have low average access time
- Storage is configured using RAID 1

Legend
Single HDD disk
Two HDD disks
Four or more HDD disks / 1-2 SSD disks
Several SSD Disks

MPEG-4/H.264/H.265 - 1-40% Motion						
Average bitrate per unit	Number of units					
	10	25	50	100	250	500
1.25 Mbit/s						
2.5 Mbit/s						
5 Mbit/s						
10 Mbit/s						

MJPEG - 1-40% Motion						
Framerate per unit	Number of units					
	10	25	50	100	250	500
1 FPS						
5 FPS						
10 FPS						
30 FPS						

MPEG-4/H.264/H.265 - 40-70% Motion						
Average bitrate per unit	Number of units					
	10	25	50	100	250	500
1.25 Mbit/s						
2.5 Mbit/s						
5 Mbit/s						
10 Mbit/s						

MJPEG - 40-70% Motion						
Framerate per unit	Number of units					
	10	25	50	100	250	500
1 FPS						
5 FPS						
10 FPS						
30 FPS						

MPEG-4/H.264/H.265 - 70-100% Motion						
Average bitrate per unit	Number of units					
	10	25	50	100	250	500
1.25 Mbit/s						
2.5 Mbit/s						
5 Mbit/s						
10 Mbit/s						

MJPEG - 70-100% Motion						
Framerate per unit	Number of units					
	10	25	50	100	250	500
1 FPS						
5 FPS						
10 FPS						
30 FPS						

Note: If enabling archiving, the recommendations in above tables can still be used by reducing the number of cameras listed in the tables to a half.

Average bitrate per unit	Number of units					
	10	25	50	100	250	500



Framerate per unit	Number of unit					
	5	12	25	50	125	250

Archiving

As with the recording database, the recommendations for using archiving depend on the needs and requirements.

In general, with many devices and long retention times, it is recommended to use archiving. Below tables provide a general recommendation for when to use archiving in various scenarios. Archiving is recommended for boxes with a checkmark.

Archive recommendation					
Number of devices	Retention time				
	1 day	7 Days	14 Days	1 Month	3 Months
10	-	-	-	✓	✓
50	-	-	✓	✓	✓
100	-	✓	✓	✓	✓
200	✓	✓	✓	✓	✓
400	✓	✓	✓	✓	✓
800	✓	✓	✓	✓	✓

Assuming the storage system for the archive is fast enough to sequentially write the needed data in the time between archives, then all disk and storage technologies as well as RAID levels are equally suitable for storing the archive database.

Reducing framerate when archiving

If using MJPEG and storing the recordings for a long period, reducing the framerate over time can be a good way to reduce the amount of storage needed.

When using MPEG-4, H.264, H.265 codecs, the reduction in needed storage is typically smaller because the keyframe in each GOP can use up to 60-80% of the data in the GOP. Furthermore, the process of archiving the recordings puts an extra load

on the recording servers and storage system, so in most cases the benefit does not outweigh the cost.

Database repair

Should it happen that the recording server is turned off without being properly shutdown, the media databases will not have been closed properly. In that case they need to be checked for issues and potentially repaired when the recording server is started again. Normally, the check and a possible light repair is quite fast and is completed within minutes.

However, if running hundreds of devices on the recording server and/or if storing recordings for several months without using archiving, the check and light repair may take longer to complete - during that time the recording server will not be recording.

To avoid such scenarios the following is recommended:

- Enable archiving and configure it to archive every ~4 hours. This will reduce the size of the recording database that needs to be checked and potentially repaired
- Additionally, configure more storage definitions (using the same disks) on the recording server, and distribute the devices across these storage definitions. This will allow multiple databases to be checked and repaired in parallel

Should the check results in some databases needing a larger and more thorough repair, the databases in question are taken offline and the recording server starts to record to the new databases.

The offline databases are then repaired one by one in the background while the recording server runs and records. Once the database repair for a device is completed, the database is put online again, and the repaired recordings for that device can be played back by the XProtect VMS clients.

Multiple storage configurations

When storing recordings for more than 20,000 'device-days' it is recommended to define two or more storage configurations in the recording server and distribute the devices across them. These storage configurations do not need to record to separate disks. They can use the same disk(s) as long as the recordings are stored in different folders.

The reason for this recommendation is that large scale and long-term testing have shown that disk performance in windows slowly degrades once a very high number of folders and subfolders exist in the path used for recording.

A 'device-day' is equal to a single device, storing recordings for one day. This means that 'device-days' can be calculated by multiplying number of devices by the retention time.

For example, 100 cameras, 100 microphones and 100 days of retention in the same storage configuration $(100 + 100) * 100 = 20,000$ 'device-days'.


If two storage configurations are used and share the same drive, each running 50 cameras and 50 microphones, the 'device-days' will drop to 10,000 per storage configuration ensuring constant high performance.

Number of recording servers

In XProtect VMS installations with thousands of cameras it may seem like a good idea to run as many cameras per recording server as possible in order to keep the number of servers and storage systems as low as possible.

However, although a single recording server can run more than 500 cameras with related audio and metadata devices enabled, it is often more cost efficient to run fewer cameras per server.

The reason for this is that the ratio between cost and performance for the servers and storage systems doesn't scale linearly. Moreover, the performance of the individual disks in the storage system doesn't scale linearly neither with more devices recording to it as the disk access becomes even more fragmented and non-sequential.



With that said, it doesn't mean that a recording server should run as few cameras as possible as that also becomes expensive and requires a lot of servers and physical space as well as maintenance.

Depending on configuration, the sweet spot between performance and cost is typically in the 200-300 cameras per recording server range.

Another advantage of not running as many cameras as possible per recording server, is that it makes it cheaper to grow the installation with more recording servers, enable use of failover recording servers or replacing broken servers.

Codec

Typical codecs used in XProtect VMS installation nowadays are: MJPEG, H.264 & H.265 - including smart variants. Assuming the same image resolution and framerate the five codecs have the following characteristics.

	MJPEG	H.264	H.264 Smart codec	H.265	H.265 Smart codec
Format	Single image	GOP	GOP	GOP	GOP
Latency	Low	Medium	Medium	Medium	Medium
Bandwidth & Storage needs	Very high	Medium-Low	Low	Low	Very Low
Suited for manually controlled PTZ cameras	Yes	Not optimal	Not optimal	Not optimal	Not optimal
Processing needed for decoding	Low	Medium	Medium	Medium-High depending on CPU/GPU	Medium-High depending on CPU/GPU

H.264/H.265 and their smart variants are normally the best choice of codec. However, in some cases like for instance with manually controlled PTZ cameras, MJPEG may provide a smaller latency and better user experience. In this case it is recommended to configure two streams for the PTZ camera. One stream using MJPEG for live viewing and PTZ control, and a second stream using H.264/H.265 for recording.

GOP length

In the XProtect VMS, the default GOP length for MPEG-4, H.264 and H.265 is 1 second. The GOP Length can be adjusted as needed. When doing so it will have an impact on the XProtect VMS compared to the standard 1 second GOP.

	Shorter GOP e.g. 0.5 second long	Standard GOP 1 second long	Longer GOP e.g. 2 second long
Bandwidth	More needed	Standard	Less needed
Storage space	More needed	Standard	Less needed
Storage performance	Lower	Standard	Higher
Video quality	Higher	Standard	Lower
Load of doing VMD on keyframes	Higher - VMD is done twice as often	Standard	Lower - VMD is done half as often
Load of doing VMD at 1 second intervals	Standard	Standard	Extremely high – cannot be done on keyframes only, so all video images must be decoded
Decode and show live video	Faster to show initial video	Standard	Slower to show initial video
Decode and display a random recorded image	Less resources needed and faster to display	Standard	More resources needed and slower to display

Windows recommendations

Since the XProtect VMS runs on a Microsoft Windows OS, there are several functions and settings in the OS on how it handles the storage system disks that impact the XProtect VMS performance. Below recommendations should be observed to obtain the best possible XProtect VMS performance.

Separate OS and XProtect VMS drives

With the recording servers it is important to use separate physical disks (different partitions on the same disk is not good enough) for the Windows OS and the XProtect VMS recordings.

In smaller installations where one or only a few individual disks are sufficient to store the XProtect VMS recordings, the drives storing the media database will fail at some point due to wear and tear caused by the constant high load of the XProtect VMS data being recorded.

If using the same drive for both the Windows OS and the XProtect VMS media database, the Windows OS installation is also lost when the disk fails. Recovering from

this scenario requires a lot of work by having to reinstall and configure both Windows and the XProtect VMS recording server.

In larger installations a storage system with RAID 5, 6 or 10 will probably be used for the drive storing the XProtect VMS media database. In such a case, if disks fail, data is not lost. However, even when using RAID, it is recommended to use separate disks for the Windows OS and the XProtect VMS media. The reason is that there is a risk that Windows at some point will perform various operations that require a lot of disk access, which potentially could impact the XProtect VMS recording performance and lead to lost recordings.

Disk formatting

To get the best disk performance it is strongly recommended that the disks are formatted with NTFS with the allocation unit size set to 64K.

Note: Do not use ReFS disk formatting as this is known to cause stability and performance issues.

HDD short stroking

When using HDDs for recording and archiving, the disk performance varies depending on where on the plates the data is stored.

The reason for this is that the tracks are longer closer to the plate's edge and thus hold more data compared to the ones closer to the center of the plate. As the plates rotate at a constant speed, the data transfer per revolution is therefore higher closer to the edge resulting in a better performance.

This behavior can be utilized to 'trade disk size for speed' by partitioning the disks to only use a smaller percentage of the disk space – for instance 60-80% of the total size. This smaller partition will be allocated space on the plates ranging from the edge in towards the center of the plates.

Now when data is written to this partition, it will always be located on the fastest 60-80% of the disk giving a higher data throughput. Furthermore, when only using a smaller percentage of the disk, the access time also becomes shorter as the disk arm and head do not need to move as far between tracks any more – thus the term 'short stroking'.

Windows search indexing

'Windows Search Indexing' can impact recording performance when it indexes the media database files. Therefore, it is recommended to either disable 'Windows Search Indexing' or change its configuration so it does not index the drives storing the XProtect VMS media database files.

Windows disk defragmenter

'Windows Disk Defragmenter' can impact recording performance when it defragments the disks. Therefore, it is recommended to configure 'Windows Disk Defragmenter' so it doesn't defragment the disks storing the XProtect VMS media database files. The XProtect VMS itself will ensure that all database files are written so they are not fragmented.

SSD trim

If using SSDs for recording it may, depending on the specific SSDs used, be necessary to schedule a 'trim' operation on the SSDs. The reason for this is that some SSDs need to be 'trimmed' from time to time to maintain optimal performance.

Usually Windows does this automatically when there is little activity on the SSD. However, with an XProtect VMS continuously recording to the SSD, there will never be a time with little activity.

As the need for trimming and the tools/methods used depend on the specific SSD and Microsoft Windows version being used, it is recommended to consult with the SSD vendor for recommendation on how to ensure that the SSD will keep running at optimal performance.

Virus scanner

The virus scanner on the XProtect VMS recording server may interfere with recording performance as it scans the data and files being written to disk.

To prevent this, certain folders and file types must be excluded from the scan.

For detailed information on how to properly configure the virus scanner, please refer to the 'XProtect VMS - Administrator manual' which can be found here:

<https://www.milestonesys.com/support/help-yourself/manuals-and-guides/?prod=3&type=13&lang=27>

Windows update and maintenance

To ensure proper XProtect VMS operation while keeping Microsoft Windows properly updated the following is recommended.

- Windows Update must be configured to download updates, but not automatically install them. Instead, a schedule must be defined for manually installing the updates at a time that minimizes impact on the XProtect VMS operation
- When the server running the recording servers requires a reboot or needs to be shutdown, or if the XProtect VMS recording server needs to be stopped for some reason, it is important to allow the recording server service time to close the media databases and stop normally by itself – even if it takes several minutes to complete. The service should never be terminated via the Windows' 'Task Manager' as it may leave media databases in a corrupted state.

Troubleshooting

Once an XProtect VMS is deployed and running, it may be needed to investigate and troubleshoot the recording server and storage performance if issues are experienced.

According to Milestone Support, performance issues often originate in the storage system not being able to cope with the load the XProtect VMS puts on it.

To investigate if the storage system of a particular system is overloaded, Windows Perfmon can be started on the server and configured to capture and save values for below counters for an extended period.

Perfmon counters	Recommendation
Processor Information(_Total)\% Processor Time	Should not exceed 80%
PhysicalDisk(<DISK>)\% Idle Time	Should be above 20%
PhysicalDisk(<DISK>)\Avg. Disk sec/Read	Below 20 ms
PhysicalDisk(<DISK>)\Avg. Disk sec/Transfer	Below 20 ms
PhysicalDisk(<DISK>)\Avg. Disk sec/Write	Below 20 ms
VideoOS Recording Server Database(_total)\Average GOP Write Time (ms)	Below 20 ms
VideoOS Recording Server Pipeline(_total)\Medias in queue	Between 0 and 100
VideoOS Recording Server Pipeline(_total)\Medias lost/sec	Should be 0 all the time

The results, either as a graph in Perfmon or as data in a .csv file format, should be examined to check that the values in general stay within the recommended range, and don't fluctuate to exceed the recommendations.

If the values indicate that the storage system is overloaded, one of the below cases might be the cause for it;

- Too many devices are recorded to the storage system
- The devices are recording with a too high bitrate
- If recording based on motion detection, more cameras detect motion at the same time than anticipated
- Too many cameras are configured to record MJPEG streams
- Archiving is enabled, causing increased load during periods when recordings are archived
- Many users play back or export recorded media simultaneously
- Third-party integrations read recorded XProtect VMS media excessively
- Broken or replaced disks in a RAID system cause performance degradation

To address the experienced issues, it is recommended to implement one or more of the below recommendations even though they may require changing the settings from the preferred settings, and might be costly to implement:

- Add one or more additional recording servers and move devices from an overloaded recording server to a new recording server until the load has normalized
- If possible, increase the performance of the storage system – for instance by using a faster RAID configuration, adding more disks, or by using short-stroking
- Add a second storage system to the recording server and distribute the devices across the two storage systems

- If archiving is not used, add a smaller and faster storage system to record the live media data, and use the existing one for storing archives
- Reduce the bitrate of cameras using MPEG
- Reduce the framerate of cameras using MJPEG or change them to use MPEG

Summary

Milestone XProtect VMS offers a unique, flexible and secure storage architecture that utilizes standard IT storage solutions to deliver reliable and high-performance recording of XProtect VMS media streams and data.

The unique archive functionality in the Milestone XProtect products' storage architecture especially, provides the option to design hybrid storage solutions that utilize the best traits of different disk technologies like; SSD for non-sequential write performance and SATA/SAS for storage capacity, resulting in affordable large-scale high-performance storage solutions for the demanding XProtect VMS installations.

Furthermore, the XProtect VMS storage architecture gives XProtect VMS designers the flexibility to design optimal storage solutions that meet the needs regardless of whether the XProtect VMS is recording just a few cameras, or several hundred cameras in a mission critical installation.

Combining flexible user access permissions, secure database handling that supports encryption and digital signing with the Evidence Lock functionality for protecting recordings against deletion, the XProtect VMS storage architecture provides an effective solution for securing sensitive and mission critical recordings.

Finally, the XProtect VMS storage architecture is not limited only to usage in the XProtect VMS recording servers. Using the XProtect Smart Client, the recorded media can be exported in the same secure, encrypted and signed media database format as used in the recording server. Doing so also provides the option to include the XProtect Smart Client – Player for playing back the exported recordings and verifying the integrity of the exported media to prove it has not been tampered with since originally recorded.

About Milestone Systems

Founded in 1998, Milestone Systems is the global industry leader in 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-in-class 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, visit **www.milestonesys.com**

Milestone Systems Headquarters, DK
Tel: +45 88 300 300

Milestone Systems US
Tel: +1 503 350 1100