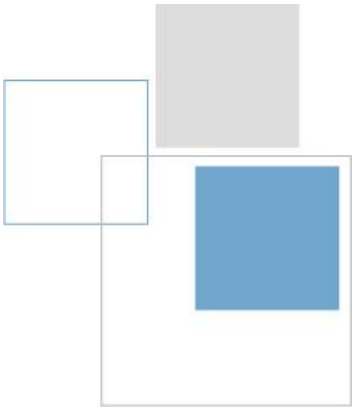

Best Practices in Video Surveillance Storage

A guide for IT and security managers
on the reliable storage of surveillance video.

milestone

Table Of Contents

Introduction

No matter what your video surveillance application, if you record video, you need to store it. Storing surveillance data preserves the crucial evidence you need to defend against an insurance/liability claim, solve a theft or violent crime, or determine the cause of a fire or explosion. In a growing number of countries and states, storage of video surveillance evidence for a set time period is actually mandated by law. As a policy, most organizations store surveillance video for at least 30 days and many for 90 days or longer.

Today many organizations are making the transition from analog (closed circuit TV or CCTV) video surveillance systems that use videotape or digital video recorders (DVRs) to Internet protocol (IP) video surveillance systems that use a number of different storage options. This raises questions about the best way to handle storage, particularly storage scalability and performance. It is becoming more important to ensure that all frames get recorded on the storage system.

This paper discusses video surveillance storage options, their evolution from DVRs and network video recorders (NVRs) to storage area networks (SANs). It also discusses the importance of selecting a storage system optimized for recording video from multiple cameras. The paper concludes by providing a technical brief on a solution using a network (SAN) based on Intransa StorStac™ Architecture and managed by Milestone XProtect™ IP video surveillance management software. In it, we provide some important best practices for the configuration of such a system.

Authors:

Eric Fullerton, Chief Sales and Marketing Officer, Milestone Systems Inc., the world's leading innovator and thought leader for open platform IP video management software.

John Dean, Director of Business Development and Alliances, Intransa, Inc. the leading innovator of shared, scalable and simple external IP storage solutions for video and IT storage.

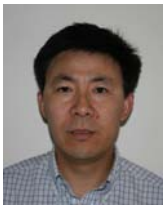

Manqing Liu, Director of Technical Marketing, Intransa, Inc.

The changing of the guard in video surveillance

Ten years ago video surveillance was primarily known as a security measure used by major corporations and financial institutions for protecting assets and monitoring perimeters. Analog cameras were strategically placed for remote viewing by security personnel, often in combination with a human guard controlling access to premises. Videocassette recorders (VCRs) recorded the captured video on tapes. Tapes were archived for a certain amount of time for post-event analysis and identification, and then reused. The recorded content was often poor because of both the limited resolution of the cameras and the use of worn-out tapes.

Then came 9/11 and, along with it, an increased desire for tougher security measures throughout the world. Fortunately, at around the same time, major technological breakthroughs in the video surveillance world ushered in significant improvements. On the heels of IP networking for computing came IP networking for video surveillance systems.

IP video surveillance management software makes it possible to centrally manage an unlimited number of networked cameras and store captured video digitally for easy retrieval and export for evidence. Using video encoders, existing analog cameras can easily be made part of the system, in addition to the newer, higher resolution (megapixel) IP network cameras that – because of a variety of advantages – are rapidly taking over the market.

The evolution of video surveillance – VCR / DVR / NVR

Recognizing the inferiority of VHS tape as a recording and storage medium, many users of analog video surveillance systems made the switch to digital video recorders (DVRs) when they became available. DVRs for surveillance applications are heavy-duty versions of the digital video recording devices popular for home use. They use hard or optical disk drives. Approximately 70 percent of the video surveillance market today uses DVRs.

DVRs serve their purpose well for small deployments, but they suffer from insufficient performance and reliability. In addition, DVRs are not scalable according to capacity. Users of these DVRs often become frustrated because of their unreliability (particularly with respect to hard drive failure), restricted storage space that limits retention periods, and the lack of a way to migrate to new IP technology, such as megapixel cameras. All of these concerns require a more scalable and robust recording solution.

Some statistics suggest that 98 percent of DVR failures are directly attributable to internal storage problems. A closer look reveals the cause. The average DVR is built with standard PC components (including hard drives designed for standard PC activity) and purchased at the least possible price from mass manufacturers. Once a DVR is put in operation, up to 16 cameras stream data continuously, 24/7. The volume of data is governed by speed of the cameras (frames per second) and resolution of the cameras (CIF rating). This results in the worst possible condition for a hard drive – a 100-percent random write-only environment. With this 100-percent duty cycle, video surveillance DVRs become drive-thrashing machines. Indeed, with 16 cameras all requiring write operations simultaneously, the result is a complex interleaving of data writing such that files are not sequential but just bursts of sequential data. Remember also that each of the camera streams have a predefined retention period. It is likely that all are the same, but it is not requirement. As retention cycles end, files (sectors, blocks, etc.) are erased, and the space is returned to the pool for reuse. In a very short time, the continued re-allocation of space results in a severely fragmented drive. This creates the worst of all possible scenarios – 100 percent random write activity. Blend in the external requests for reading previously recorded video and the randomness of the read and write accesses becomes overwhelming. It is no wonder that the internal drive is the single most failure-prone component reported by DVR owners.

These drive failures are a real problem since video surveillance systems depend on their storage system not to fail, shut down, or otherwise stop functioning. If a drive in a DVR fails, complete banks of cameras will cease to function. Organizations can ill afford these failures, whether their needs for surveillance include regulatory compliance audits; internally established standards reviews; avoidance of personal injury claims; fear of terrorist activities, theft, vandalism

and other crime; video-recorded industrial process reviews; or the reduction of property insurance rates.

There are other challenges as well with a system depending solely on DVRs for storage. One of those is the retention period itself. How does one increase the retention period to meet new regulatory requirements when there is a limited amount of storage and you want to keep all the recording settings for the cameras as defined?

Another challenge is the addition of new technology, such as IP megapixel cameras. The newer higher resolution cameras provide better quality recorded images for identification of people, objects and events, but the output from these cameras requires much more storage. For systems limited by the fixed capacity of DVRs, the best solution is to add external scalable storage.

A better choice than DVRs is the latest generation of NVRs. These NVRs are open platform, so they offer greater scalability and flexibility. (Open platform devices use Application Programming Interfaces – APIs – to allow third party solution providers to integrate with the platform to add functionality). Today's NVRs enable you to choose from the best IP video surveillance management software products available and the best storage choices. Coupled with excellent IP video surveillance management software and external scalable storage, NVRs can support a much higher number of cameras (including megapixel cameras) than traditional NVRs and DVRs.

Table 1: Comparison of DVRs and Today's NVRs

DVR	Modern NVR
• Dedicated box with analog video inputs and internal storage • All intelligence at DVR • Resolution limited by what's available from analog cameras • Most have finite number of inputs: 4, 8, 16, 32 • Finite frame rate capacity: 30, 120, 240, 480 • Limited storage capacity	• 100 percent IP-based video • Handles input from IP cameras, video servers, and analog cameras equipped with digital encoders • Enables intelligence at the camera level as well as NVR level • Enables high resolution megapixel cameras • Enables remote viewing and recording

Scalable storage for video surveillance

With the increasing popularity of IP video surveillance systems, data storage in video surveillance is undergoing a major revolution. Today's video surveillance systems require storage that can:

- Scale to larger storage capacities
- Record higher frame rates without dropping frames
- Add cost-effective storage for longer retention periods
- Handle higher resolution video from megapixel and multi-megapixel cameras
- Store and manage videos centrally from a distributed implementation
- Scale to accommodate future growth

Increasing storage needs have led many users of IP video surveillance systems and hybrid systems incorporating both analog and IP network cameras to consider storage area network (SAN) for their storage requirement. A SAN is a dedicated network, separate from LANs and WANs, that is generally used to connect numerous storage resources, such as DVRs and NVRs, to one or more centralized, shared storage arrays. SAN arrays provide faster block-level (as opposed to file-level) access to storage – this means a network host (such as an NVR) views a remote storage server as if it were a locally attached drive and transfers information in blocks as if it were accessing a local hard drive.

Fibre Channel (FC) technology is currently the dominant infrastructure for storage area networks (SANs). The Fibre Channel protocol and interconnect technology grew from the need to provide high performance transfers of block data. However, FC SANs are very expensive. In addition, they require a special skill set to manage. Gaining this skill set involves expensive specialized training.

Internet SCSI (iSCSI) is an industry standard developed to enable transmission of SCSI block commands over the existing IP network by using the TCP/IP protocol. iSCSI is a technological breakthrough that offers organizations the possibility of delivering both messaging traffic and block-based storage over existing Internet Protocol (IP) networks, without installing a separate FC network.

IP SAN virtualizes hundreds or even thousands of physical disk drives and presents a logical drive to an NVR through the operating system in the form of a logical unit number (LUN). This is represented via a drive letter. Physically, these drives form a disk group upon which a data protection technology known as Redundant Array of Independent Disks (RAID) is constructed. There are different levels of RAID offering different levels of redundancy, from none at all (RAID 0, striping) to completely mirrored data (RAID 1), to various parity constructions (RAID 3, 4, 5, 6) where the loss of one hard drive doesn't affect the integrity of your data. (RAID settings and their effect on capacity utilization will be discussed more fully later in the paper.)

Typically a SAN storage array provides increased availability, resiliency and maintainability by using additional, redundant components (controllers, power supplies, fans, etc.) in an attempt to eliminate all single points of failure (SPOFs). Additionally, these components are hot-swappable. IP SANs, so constructed, solve the three primary challenges faced by traditional DVRs: reliability, retention and resolution. They also offer a scalable storage platform for NVR deployment. A SAN, such as an Intransa IP SAN, can enhance video storage performance through video surveillance workload optimization.

Connecting to an iSCSI SAN storage array is a simple matter. Many NVRs already have Network Interface Cards (NICs) available for local area network (LAN) connectivity. In addition, it's very easy to install another NIC if desired to attach an iSCSI SAN storage array. NVRs connect to the iSCSI SAN storage array through the IP network. From an NVR's perspective, the iSCSI SAN is presented as another drive letter. The DVR/NVR application, once configured to point to the new drive letter, will do the rest. Figure 1 below provides a diagram of how this works.

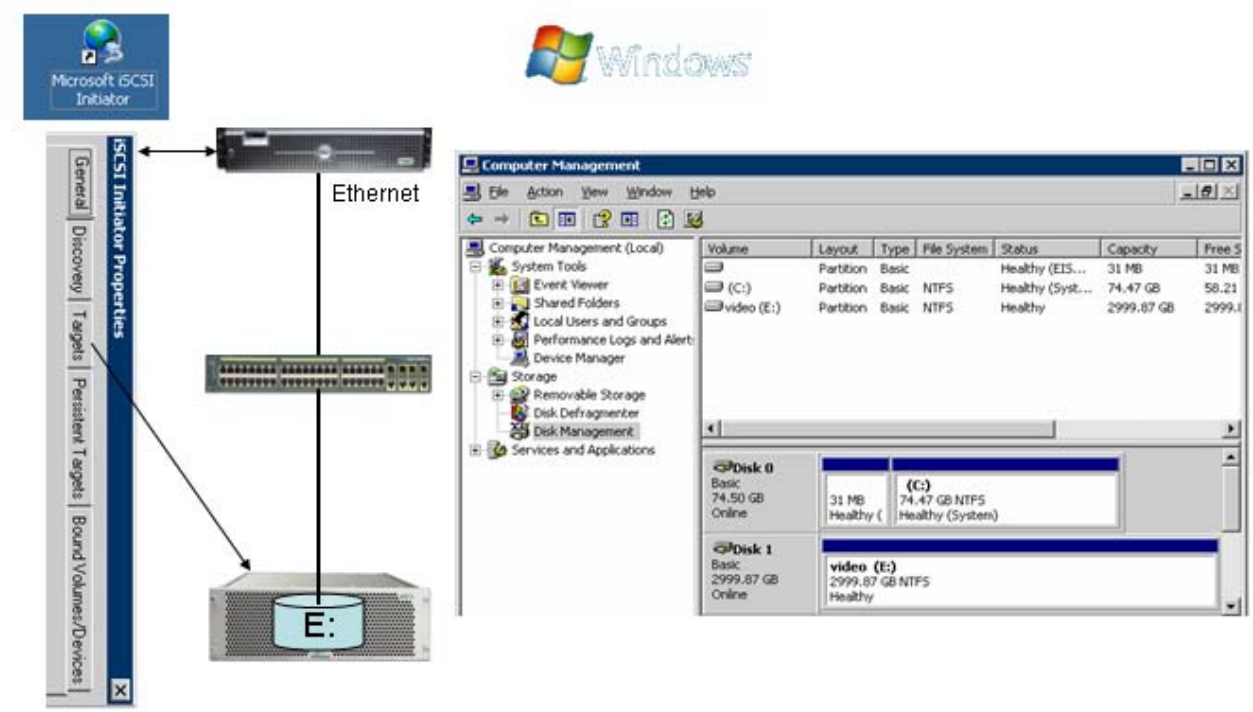

Figure 1. Typical iSCSI SAN Storage Array Representation

The place to start for any storage system? Selecting the right software.

The heart of any IP video surveillance solution is its management software. This software enables you to manage and control everything from your cameras to your video storage. It's important from the start to choose an IP video surveillance management software solution that future proofs your system by allowing you to design the system to fit the current and future needs of your organization's security goals. This can include everything from providing support for unlimited cameras to interfacing with the widest choice of video storage solutions.

In particular, you'll want software that can make effective use of existing IP networks through all the various video compression techniques (MJPEG, MPEG4, MPEG4 ASP*, H.264* and MxPEG). You'll also want to be sure the software uses bandwidth-optimized multi-streaming to split a single video stream from a camera into differentiated streams for simultaneous live view and recording. With software like Milestone XProtect a user can request a live view at a different frame rate and resolution than the recording settings. This ensures you get the video recorded for storage at the selected frame rate and resolution no matter how it is being viewed.

Another important feature to look for in video surveillance management software is the ability to optimize the value of storage by enabling use of iSCSI SAN technology with SCSI drives for short-term recording and SATA II drives for mass-scale, long-term archiving. Milestone XProtect, in fact, is unique in its ability to move historical video data multiple times per day to network drives while maintaining all recording functions. XProtect can also enable you to configure individual scheduling and retention time per camera or camera group. Hourly to daily database archiving, with the option to automatically move saved video to a network storage solution, conserves storage capacity on the local server, yet keeps video available for easy playback. Naturally, the more capable your SAN with respect to the needs of video surveillance storage, the better the performance you'll achieve with this special capability of XProtect.

The IP video surveillance management solution you choose is the software you will rely on to integrate all your system's elements, including cameras, DVRs, NVRs, and SAN. It's also your interface and control center for operating them all. Consequently, the first place to start in designing a new IP video surveillance system or transitioning to a hybrid solution that uses both analog and IP network cameras is choosing a full-featured IP video surveillance management solution. You want an open platform that gives you the greatest flexibility in selecting the other components of your system, setting up your SAN, and scaling to meet future needs.

Achieving greater reliability through dual-stage archiving

A unique capability of Milestone Systems XProtect IP video surveillance management software is its ability through scheduling to enable a SAN to do dual-stage archiving. This is an important capability from a performance, reliability and cost point of view. Historically, dual-stage archiving enables you to record primary “day one” storage on faster, more reliable (and generally more expensive) hard drives and schedule daily transfers of the recorded video onto less expensive hard drives for long-term storage.

SCSI drives are specifically designed for the intensive random write-to-disk needs of applications like video surveillance. They provide:

- 15K RPM rotational speeds (versus 7200 RPMs for most SATA drives) for significantly lower latency, greater I/O processing capability, and up to 50% greater sustained throughput per drive than a typical SATA drive
- Greater ability to withstand higher rotational vibration interference
- Lower mean-time-between-failure (MTBF)
- Lower bit error rates (BER)
- Warranties of five or more years

These features and capabilities make SCSI drives ideal for providing day-one storage for SAN installations serving large numbers of cameras and environments where reliability and performance are more important than cost and capacity. SATA drives make more sense for serving smaller numbers of cameras, recording on event, and long-term storage applications employing RAID to ensure adequate redundancy.

SATA or SCSI? The best practice for SANs for IP video surveillance storage today is to either use both or SATA II systems like the Intransa StorStac™ Architecture that are specially configured for optimal performance and reliability. In the past, there wasn’t the option to use SATA and SCSI drives in the same server because SCSI drivers weren’t serial like SATA drives. But Serial Attached SCSI (SAS) disks have changed that and now the same serial controller can control both SCSI and SATA drives. This significantly reduces the cost of using both types of drives, plus decreases training and support costs for a lower cost of ownership. You no longer have to have special purpose hardware. It’s all “off the shelf.” What’s more, scalability is excellent. You can add a SCSI or SATA drive, or both, to the array and reconfigure on the fly. In such hybrid storage systems, SCSI drives, because of their superior availability and reliability, are used for day-one storage and SATA drives (employing RAID to ensure adequate redundancy) are used for long-term storage. In storage systems like the Intransa solutions, similar results are achieved using SATA II drives for both primary and secondary storage.

With the Intransa iSCSI SAN storage system, XProtect IP video surveillance management software can execute dual-stage archiving on the same storage array, achieving both performance and archiving at an optimal cost of ownership. What's more, the user of XProtect doesn't need to be worry about matching storage technologies to functions.

In this paper's Technical Brief, configuration guidelines for optimizing the Intransa iSCSI SAN storage system for performance will be provided. Intransa's iSCSI SAN storage system is totally constructed with SATA II technology drives.

The advantages of SATA II

A new specification in hard drive technology, SATA II, makes SATA more appropriate for enterprise environments and for long-term storage (and even primary day one storage) in video surveillance applications. SATA II adds three important features: port multipliers, port selectors, and native command queuing.

- **Port multipliers:** The port multiplier specification from the Serial ATA Working Group allows up to 15 drives to be connected to a SATA controller via a port multiplier. This makes it much easier to build disk enclosures using SATA II drives. (With SATA I, parallel ATA drives had to be configured as master and slave, and daisy-chained from each controller. This made it harder to scale.)
- **Port selectors:** This feature allows two hosts to be connected to one drive. This is useful because it creates a redundant connection to the disk drive. If one of the hosts has a failure, the second host, acting as a spare, can take over, so that access to the storage is maintained. This type of redundancy is essential for enterprise environments.
- **Native command queuing:** This feature improves the performance and efficiency of SATA II drives. Normally commands will arrive at a disk to read or write from different locations on the disk. With SATA I, commands are executed in the order they arrive, creating a great deal of mechanical overhead because the read/write head is constantly being repositioned. SATA II drives use an algorithm to determine the most efficient order to execute commands. This both reduces mechanical overhead and improves performance.

While not as fast or tough as the more expensive SCSI drives, SATA II drives provide a strong alternative, particularly in storage solutions that have been specifically optimized through their storage controllers and other components for video surveillance.

Technical brief: An example of a SAN from Intransa

Intransa offers proven IP storage solutions that integrate with existing infrastructures, including analog and digital cameras, DVRs, and NVRs. Intransa scalable video surveillance storage solutions that work with Milestone IP video surveillance management software and many other video surveillance management, monitoring and streaming applications without modification. Intransa IP storage can be grown modularly in a pay-as-you-grow model to support hundreds, thousands or tens of thousands of cameras, all over standard IP and Ethernet, without ever running out of ability to grow.

The Intransa StorStac™ Architecture system consists of two independent scaling components:

- Performance Controller Units (PCU), which handle all storage management and virtualization
- Storage Capacity Enclosures (SCE), which are the disk drives and the disk drive enclosures

PCUs are clustered as a single realm (addressable storage) to manage the performance scaling and high availability requirements of video surveillance. Such independent scalability allows administrators to manage the number of cameras/resolution independent of the retention period. The figure below is a depiction of the system.

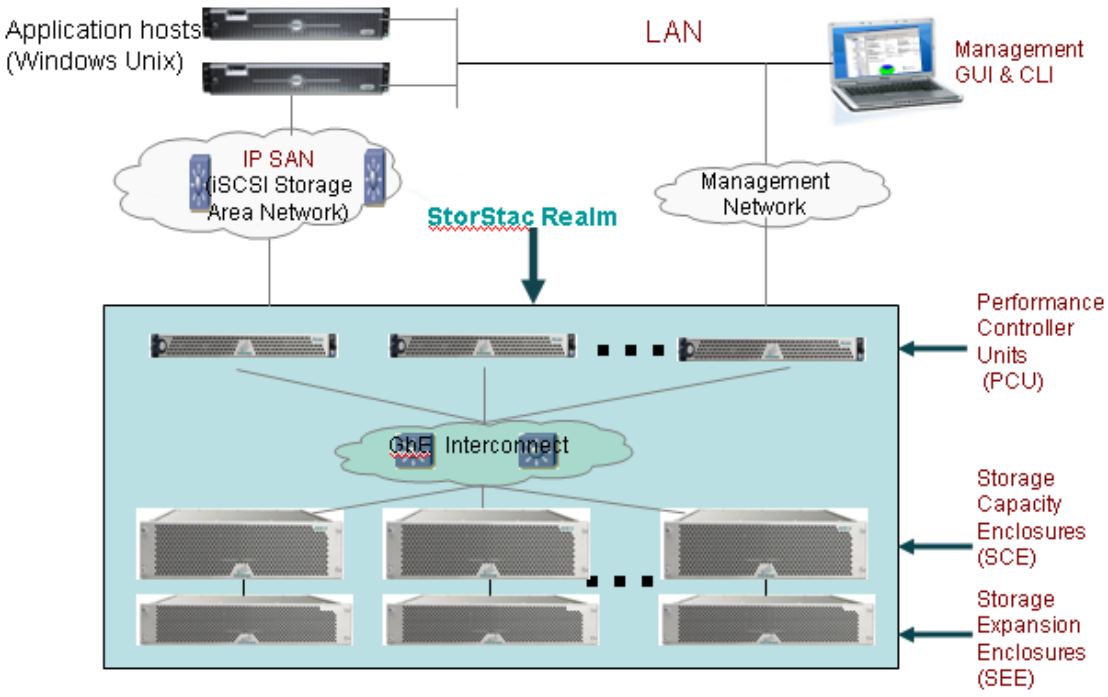

Figure 2. Intransa BuildingBlock IP SAN Storage System

As depicted in Figure 2, multiple PCUs are clustered for load balancing and complete failover support. Up to eight of these PCUs can form a cluster. Connection can be accomplished with 1 GbE or 10 GbE interfaces. With the former, throughput for the system can range from 220 MBps (megabytes per second) to 880 MBps. With the latter, system throughput begins at 700MBps and can exceed 5000MBps.

Connectivity to the storage is managed through back-end interconnects and all connections with the Storage Capacity Enclosures (SCE) are IP-based. Each of the SCEs is capable of supporting up to three Storage Expansion Enclosures (SEEs).

This architecture facilitates scalable storage from as small as 4TB to 1,500 TB and beyond. The entire system is easily managed using a single IP address via a graphical user interface (GUI) that has been designed to eliminate the need for the video surveillance administrator (VSA) to understand storage nomenclature and instead concentrate on the real issues of camera type, resolution, frames per second, compression algorithm selection, and retention periods required for each. An example of the GUI follows.

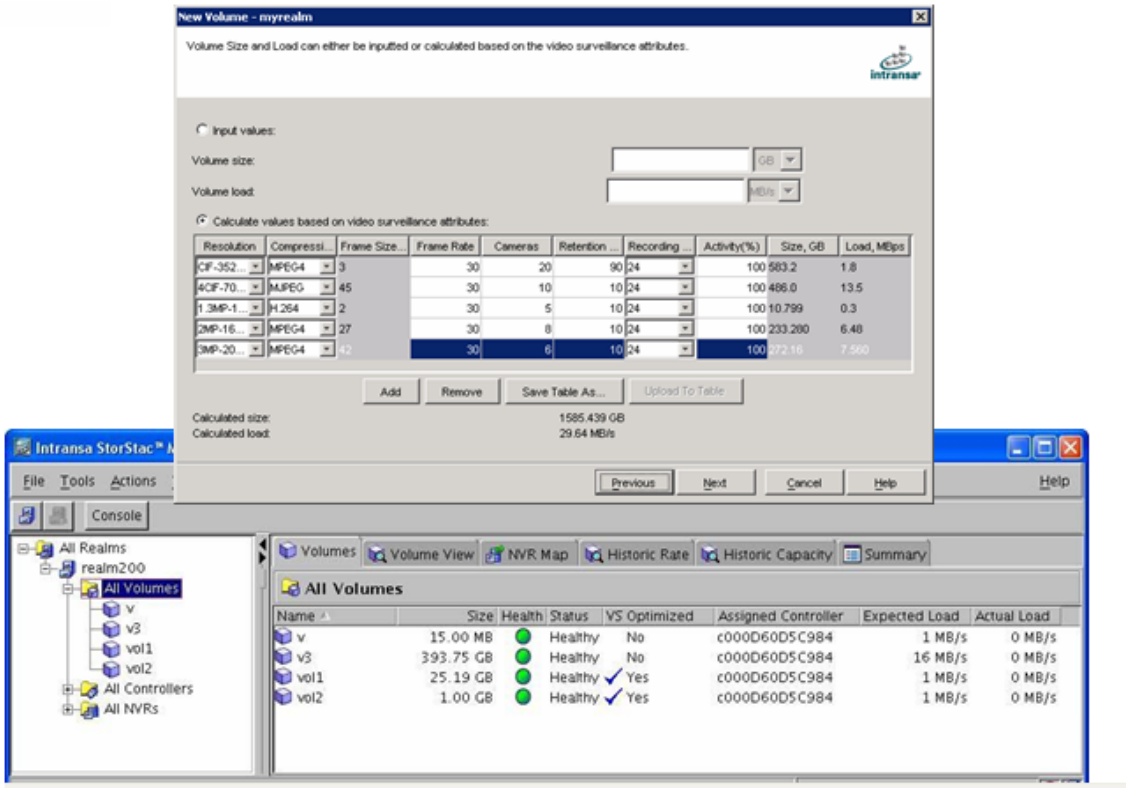

Figure 3. Sample of the VSA GUI

The Intransa management system graphically informs the administrator of all pertinent statuses of the system, including: analysis of operations, activity from channel streams, alerts and thresholds. Since the system is IP-based, no extra training or expensive equipment, such as Fibre Channel cards or switches, is required. The complete operation can be run by the same IT team that runs the rest of your IP networking operations and integrated into an IP video surveillance management software solution like Milestone XProtect.

The Intransa IP SAN scalable storage system is a SATA II drive-based system – a technology that is proven, rugged and green. SATA II systems are supplied with drives having 500 GB, 750 GB and 1,000 GB (1 TB) capacities in 3.5-inch dual ported containers supporting 7500 rpm. This technology is precisely suited for streaming media applications such as video surveillance. This technology is designed for large record format, high density storage, and is sufficiently rugged to sustain the heavy random write operations from a video surveillance workload. In addition, energy consumption is 50 percent of drives using Fibre Channel technology. With the lowest cost per GB available in the market, SATA II provides the best total cost of ownership and is the correct product for video surveillance applications.

Intransa IP SAN scalable storage solutions scale from 4 TB to more than 1,500 TB and from 200 MBs throughput to more than 3,000 MBs throughput on the performance curve. The end result is that this IP SAN scalable storage solution is the correct upgrade to solve existing limitations and facilitate the addition of new technology. It's also ideal for new video surveillance installations. It grows with your video surveillance operation as the need increases for more storage and greater performance.

Best practices in configuration of a video surveillance storage system

Configuring a SAN storage system for optimal performance requires a systems approach. The DVRs/NVRs, the network, and the iSCSI SAN scalable storage array all need to be optimally configured for making the best use of the entire system. Intransa and Milestone have spent a significant amount of time understanding these parameters and document a few key ones here.

There are several elements involved in the configuration of a video surveillance storage system that deserve better scrutiny. Clearly, in video surveillance, streaming video data to spinning disk storage requires some very well thought-out configuration parameters. These parameters are essentially best practices for each of the operational components: the video surveillance host DVR/VNR, the networking switch/router, and the iSCSI storage platform. In the following several pages, we will address some of the more directly controlled operations.

Hardware Preparations

For NVR hardware, you will need to make sure you have enough resources (such as CPU, memory and NIC bandwidth) to accommodate both the number of cameras supported and the speed at which they record. Talk to your NVR vendors to find out the proper hardware recommendations for CPU and memory. For Milestone Video Management Software, a typical Dell 2950 has sufficient resources to support over 2000FPS at 4CIF MJPEG resolution.

If there is no Gigabit network adapter card already present, install one on each DVR/NVR. If high availability is your concern, consider using dual-port NICs or two NICs and two switches in redundant configuration. Figure 4 shows a typical deployment diagram.

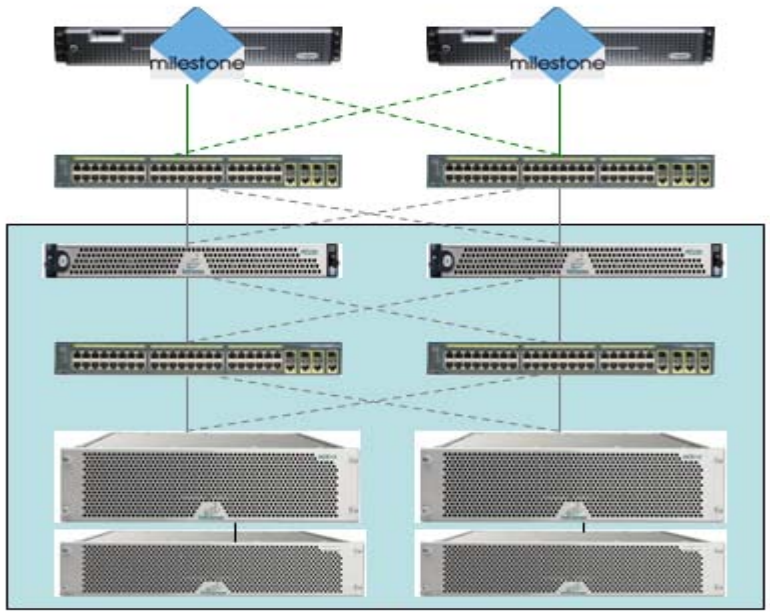

Figure 4. Typical iSCSI SAN Deployment

iSCSI initiator driver installation on NVR

The Microsoft iSCSI initiator service enables the host server to connect to the iSCSI volumes on the storage array. iSCSI initiator is freely available from Microsoft. It is available on all Windows platforms, including Windows XP/Vista, and Windows 2000/2003/2008 server. For Windows XP/2000/2003, you will need to download and install it if it has not been installed. Select the MPIO option during installation which will enable the NVR to have multiple paths to the storage system for path redundancy. Two 1GbE ports from Intel PRO1000 NIC are used for the iSCSI connection for redundant configuration by taking advantage of MPIO.

You can find the Microsoft iSCSI Software Initiator (http://go.microsoft.com/fwlink/?LinkId=28169) on the Microsoft Download Center.

MPIO not only offers high availability but also increases the performance. If one link between the NVR and storage system is down, IOs will be redirected to the second link automatically. MPIO is available for Windows server 2000/2003/2008.

Host NVR best practices

For a typical NVR, you need to ensure you have sufficient resources to manage the complete application – CPU, memory, and network interface card (NIC) bandwidth. Since the storage access is through regular NICs using iSCSI initiator, we recommend using a 1GbE-based NIC. For compatibility between iSCSI initiator and the storage, there are many available NICs in the market. Do not shortchange this element of the equation.

In setting up your DVR/NVR NICs, enable the Jumbo frame setting in the configuration setup section of the initialization mode. (Jumbo frames are Ethernet frames containing more than 1,500 bytes of payload. Formerly the maximum transmission unit was 1,500 bytes, but today's Gigabit Ethernet switches and Gigabit Ethernet network interface cards support Jumbo frames which can carry up to 9,000 bytes of payload.) Enabling the Jumbo frame setting will reduce the CPU utilization on your DVR/NVRs and help improve performance.

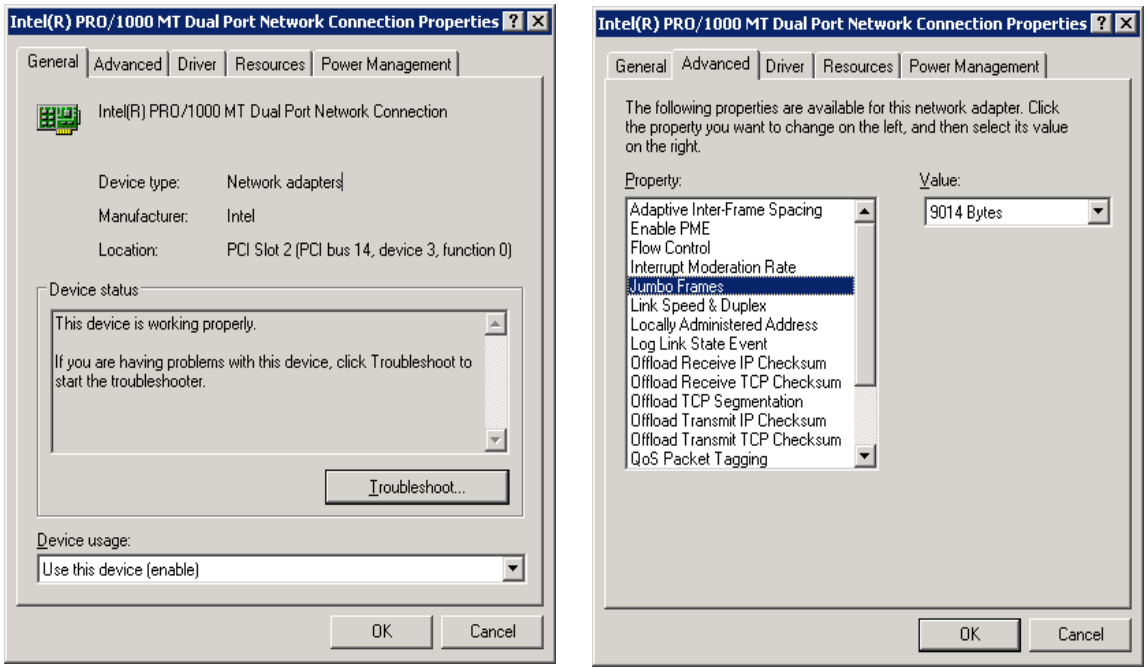

Figure 5: NVR NIC configuration

By default, most NICs do have TCP offload capabilities. TCP offloading can significantly reduce your NVR CPU consumption. On Dell 2950 which has dual 3.0GHZ CPU, Table 2 shows a typical CPU utilization for various request size. The measurement is based on 10Gb network and can also be applied to 1Gb network from CPU utilization point of view. Given that more than 90% of the requests are write IOs, the CPU overhead due to software iSCSI is very minimal.

	READ		WRITE	
Request Size (KB)	Throughput (MBps)	Host CPU Utilization %	Throughput (MBps)	Host CPU Utilization %
8	130	25	105	10
64	460	26	390	14
256	650	27	485	12
512	700	26	520	12

Table 2: CPU utilization due to iSCSI overhead on NVR, measured using Dell 2950.

Typically for 100 cameras with 30FPS at 4CIF with MJPEG compression, you will see about 10% CPU resources due to software iSCSI traffic.

You can also set up the redundant network links between DVR/NVRs and the iSCSI storage system. This is typically done through multipath IOs (MPIOs). If one link between DVR/NVR and storage system is down, IOs will be redirected to the second link automatically. MPIOs not only offer high availability, but also increase performance.

Networking best practices

We recommend using separate switches for the front-end and back-end networks. The front-end network is the network between DVR/NVRs and PCUs. The back-end network is internal to the iSCSI SAN storage system, connecting the PCUs with the SCEs.

We recommend the following configuration settings:

1. Enable Jumbo frame on DVR/NVR (as described in the previous section), the front end network, and the iSCSI SAN storage system.
2. Enable Flow Control on the switch.

Storage setup best practices

When setting up a disk in a SAN, we suggest the executing the following.

Set RAID Level: Much information is available on RAID settings and we won't repeat it here. For best capacity utilization in video surveillance storage system, we recommend RAID 5 and RAID 6. RAID 5 and 6 use data striping – the segmentation of logically sequential data (such as a single file) to enable segments to be assigned to multiple drives – and parity for data protection. This enables faster disk writing and reading. RAID 5 uses distributed parity so that three or more disks can protect against the loss of any one disk. This reduces the usable storage capacity of the array by one disk, but enables greater reliability. RAID 6 uses dual distributed parity to enable recovery from the loss of two disks.

Use GPT instead of MBR: When you first initialize a disk in your SAN, choose GUID Partition Table (GPT) instead of Master Boot Record (MBR). GPT is a standard for the layout of the partition table on a physical hard disk. Video surveillance, as you know, consumes many TBs of storage. Choosing GPT over MBR enables volumes greater than 2 TB.

Use Disk Management (accessible from the Windows Computer Management function) to convert the disk from MBR to GPT:

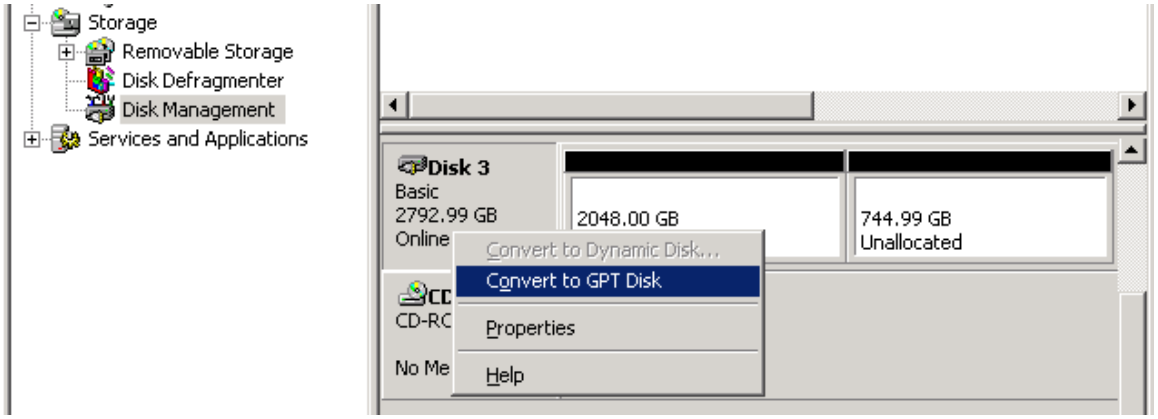

Now you will see a single volume as:

Do an alignment before you create NTFS: Here is an easy way to increase disk performance using the Command Line Interface (CLI) commands below:

C:\>**diskpart**

Microsoft DiskPart version 5.2.3790.1830

Copyright (C) 1999-2001 Microsoft Corporation.

On computer: MKT-10G-S1

DISKPART> **select disk 3**

Disk 1 is now the selected disk.

DISKPART> **create partition primary align=64**

DiskPart succeeded in creating the specified partition.

DISKPART> **exit**

Leaving DiskPart...

Format NTFS using 64K allocation unit size:

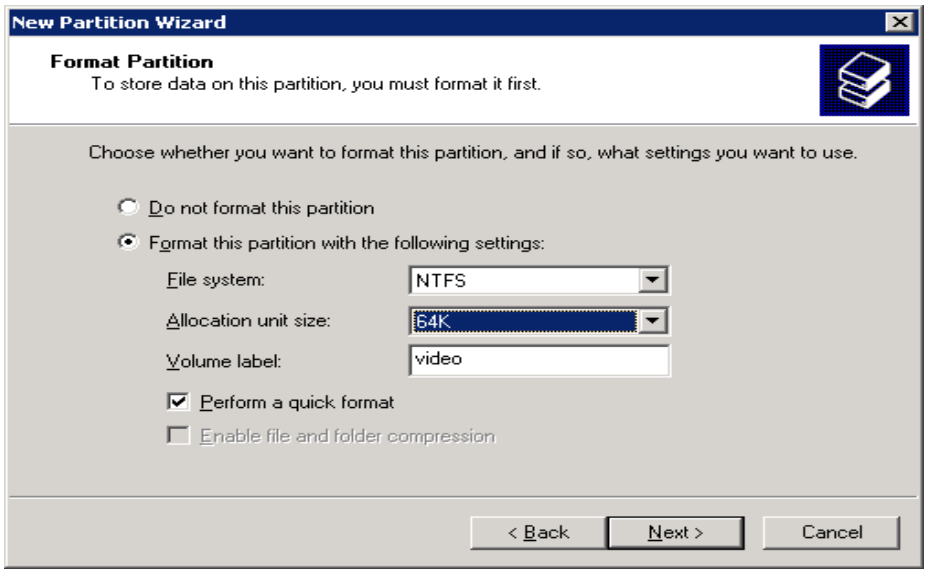

Now configure your DVR/NVR to use this storage device. For most DVR/NVRs, simply associate the NTFS drive letter (E:\ drive in this example) with a specific camera(s):

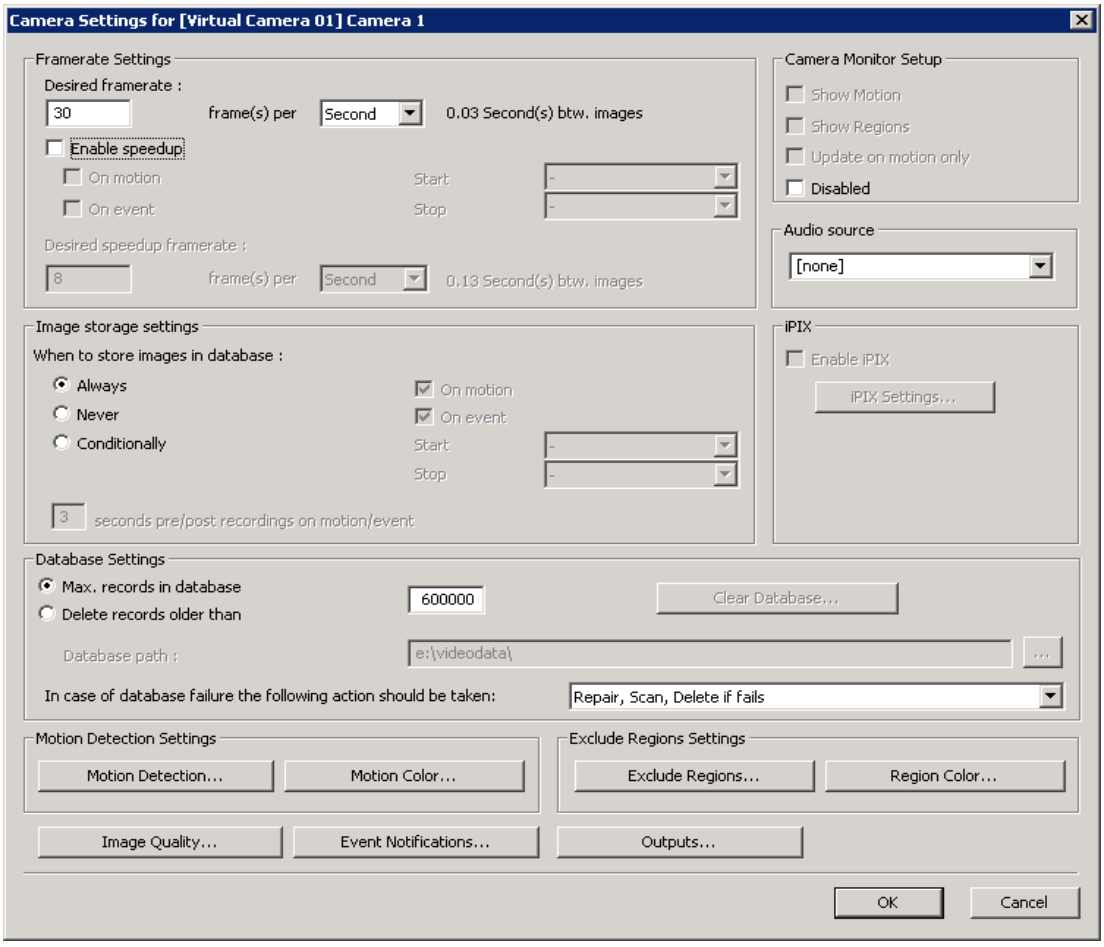

Disk Failure and RAID Rebuild: The practices we have chosen to highlight here focus on the impact of the speed of a RAID Rebuild on the ability of the system to record video streams with minimal frame loss. RAID 5 protects against data loss in the event of a drive failure. When a disk fails, a spare drive will be automatically allocated and the RAID group will begin to repair itself. The system will operate in a degraded mode during this rebuild process.

During the rebuild, it is important to have control over the rebuild speed. In some situations, it is desirable to have the RAID group rebuild as soon as possible. In other situations, the RAID should be rebuilt at a slower speed so no significant impact occurs to the real time recorded video session. The speed of the rebuild and the importance placed on avoiding frame loss are a decision you'll have to make based on the importance of the video quality in your surveillance application. Generally, some frame loss is acceptable and often unnoticeable, while excessive frame loss can hurt the value of the video as evidence.

Increasing performance through proper LUN layout: A Logical Unit Number or LUN is the term for an array of disks or drive volume that is allocated to perform and respond as a single storage device.

Proper LUN layout can increase your performance. As a matter of fact, LUN layout is very important for objectives such as high availability and reducing performance impacts due to storage controller failure, RAID rebuild due to disk failure, etc.

In a configuration with fourteen 750 GB SATA II drives, Table 3 summarizes the measured performance results with various of LUN layout strategies together with some of the pros and cons:

Disk Group	Volumes (LUNs)	Number of Cameras [30FPS, 4CIF, MJPEG] Supported	Comments
One 14 disk-RAID5	1 LUN	>50	This is the baseline.
One 14 disk-RAID5	2 LUNs	>64	With 2 LUNs using a single 14 disk RAID5 disk group, spindle resources are better utilized.
Two 7 disk-RAID5	4 LUNs	>72	With two disk groups, IOs can be distributed between two DPUs so the system can handle more cameras with the 14 drives.

Table 3: LUN Layout and Performance Optimization

Summary

This paper has introduced you to the advantages of adding a SAN to a video surveillance system currently using DVRs and NVRs as its primary storage system. It has explained:

- The reasons DVRs and NVRs with internal storage cannot keep up with the storage demands of today's IP network video surveillance systems
- How SANs provide scalable, reliable, high-speed IP storage
- How RAID types and settings can affect capacity utilization and redundancy
- Why it is important to buy a SAN designed and optimized for video surveillance usage
- Why iSCSI SANs using SATA II technology provide optimal performance, reliability and cost
- How the right IP video surveillance management solution can:
    a. Future proof your surveillance system by providing a scalable, open platform that enables you to use variety of products (including cameras and storage solutions)
    b. Optimize your system for bandwidth, performance and storage
    c. Split a single video stream from a camera into differentiated streams for simultaneous live view and recording
- The advantages of the unique dual-stage archiving capability of Milestone XProtect software that enables the use of fast, ultra-reliable local disks for short-term recording and more affordable drives for mass-scale, long-term archiving

We've also provided a technical case study of a best-in-class IP storage system from Intransa Inc. that can be grown modularly in a pay-as-you-grow model to support hundreds, thousands or tens of thousands of cameras, all over standard IP and Ethernet, without ever running out of ability to grow. In this case study, we've discussed many of the best practices in configuring such a system to maximize performance and reliability.

Much more information can be obtained from the respective companies responsible for this paper.

- Milestone Systems (www.milestonesys.com)
- Intransa Inc. (www.intransa.com)

Milestone Systems

Innovator. Milestone Systems is internationally recognized as an innovator and thought leader in open platform IP video management software. Milestone's XProtect products operate as the core of surveillance systems: connecting, sharing and managing all devices through a single interface that is easy to learn and operate.

Easy to use. The XProtect platform is easy to use, proven in operation and scales to support unlimited devices. XProtect products support the widest choice of network video hardware and are designed with an Application Programming Interface (API) that integrates seamlessly with other manufacturers' systems.

Best-of-breed. Using XProtect, you can build scalable, "best of breed" solutions to reduce cost, optimize processes, protect assets and ultimately increase value in a company's products and services.

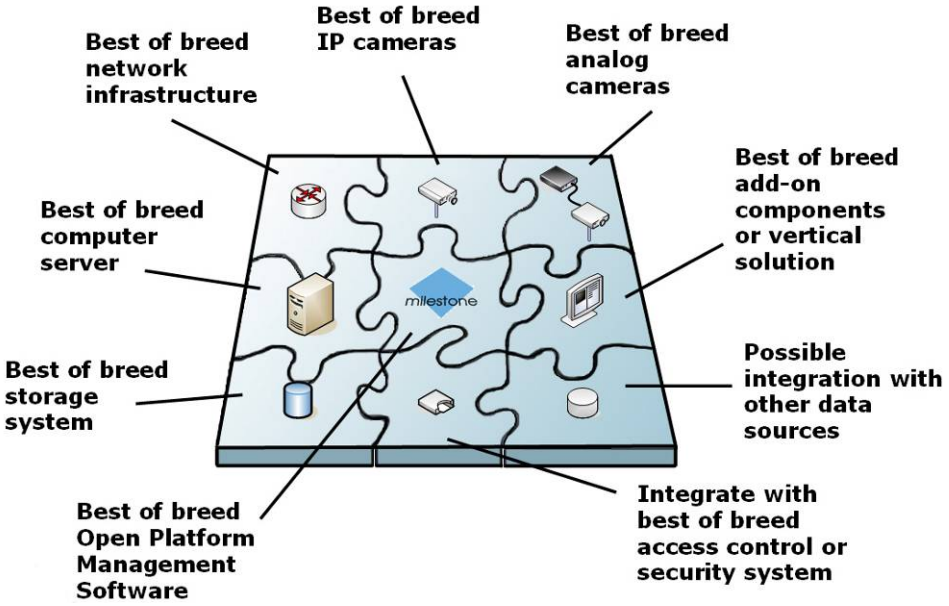

Intransa, Inc.

Intransa, Inc. is the leading innovator of shared, scalable and simple external IP storage solutions for video and IT storage. Based in San Jose, CA, Intransa offers scalable, high performance security-grade storage with outstanding availability and price/performance, with built in Set-and-Forget Management for proven ease of use.

Beginning as a 3Com spin-off, Intransa has been shipping shared IP network storage solutions since 2003. Intransa IP storage is available worldwide through authorized StorPartner security and systems integrators. Intransa is an Axis Communications Applications Development Partner, Cisco Ecosystem Provider, Microsoft Gold Certified Partner, and VMware Technology Alliance Partner along with many other industry partnerships through our StorAlliance Technology Partner program.

For more information about Intransa and our industry-leading sharable, scalable and simple external IP storage solutions, to locate a StorPartner security and systems integrator, or to join the StorAlliance Technology Partner program, please visit us at www.intransa.com.

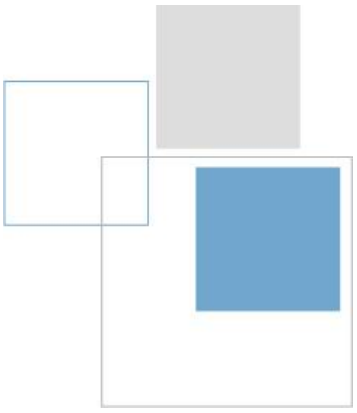

Milestone Systems is the industry leader in developing true open platform IP video management software. The XProtect™ platform gives users a powerful surveillance solution that is easy to manage, reliable and proven in more than 35,000 customer installations worldwide.

With support for the industry's widest choice in network hardware and integration with other systems, XProtect provides best-of-breed solutions to "video enable" organizations – reducing costs, optimizing processes, and protecting assets.

Milestone software is sold through authorized partners in approximately 90 countries.

Office Locations:

Milestone Systems Inc.
8905 SW Nimbus Avenue, Beaverton, OR 97008, United States
Tel: +1 (503) 350 1100

Milestone Systems A/S (Headquarters)
Banemarksvej 50, 2605 Brøndby, Denmark
Tel: +45 88 300 300

Milestone Systems DE
Am Kleefeld 6a, D-83527 Haag i.OB., Germany
Tel: +49 (0) 8072 442173

Milestone Systems Italy
Via Paisiello,110, 20092 Cinisello Balsamo, Milano, Italy
Tel: +39 02 6179 508

Milestone Systems UK, Ltd.
118 Codnor Gate, Ripley, Derbyshire DE5 9QW, Great Britain
Tel: +44 (0) 1773 570 709

Milestone Systems France
121 rue d'Aguesseau, 92100 Boulogne-Billancourt, France
Tel: +33 141 03 14 82

Milestone Systems Japan
c/o Royal Danish Embassy, 29-6, Sarugaku-cho, Shibuya-ku, Tokyo 150-0033, Japan
Tel: +81 (0)3 3780 8749

Milestone Systems Pte. Ltd.
30 Robinson Road, 13-03 Robinson towers, Singapore 048456
Tel: +65 6225 2686

Milestone Systems Middle East
P.O, Box 500809, DIC, Building 5 IEB, 6th floor Office 606, Dubai, United Arab Emirates
Tel: +971 50 8827093

Corporate web site: www.milestonesys.com

Milestone User Group & Discussion Forum www.milestonesys.com/usersgroup