The Curious Case of Video Surveillance

A Guide for IT Professionals to the World of Physical Security

V1.2

Intransa VideoAppliance™ was named 2009 New Product of the Year at ASIS Anaheim

Table of Contents

Introduction

IT professionals are often surprised to find how widely video surveillance applications commonly deployed by physical security practitioners differ from IT environments, despite the common use of networking, servers, workstations and storage that share a common heritage.

Video surveillance systems are real-time, resource intensive applications that typically run 7x24x365 in a streaming mode, where the processing power, network bandwidth, disk throughput, and storage capacity requirements are extreme. A proper design and integration of servers, networks, and storage are key to successful video surveillance systems, particularly since most systems run unattended without updates or downtime for maintenance.

For IT practitioners, thinking of video surveillance as a type of 911 system can be helpful to understanding the environment. While hoping it's most critical features are never needed, video surveillance needs to be always ready. Consider that a bad day in physical security is when someone is injured or killed, and is especially bad if the video system capturing the event has failed. This in itself is unlike many day-to-day IT applications where downtime can mean inconvenience but not result in a physical safety impact. Such is the case in physical security and the use of video surveillance.

As a result, architecting and implementing a reliable surveillance system without excessive cost and ongoing complexity is important to mission-critical video surveillance applications.

This document is intended for an IT knowledgeable audience, seeking to understand the video surveillance market and its implications to them in terms of server and storage requirements.

Background

This application originated from analog CCTV (Closed Circuit TV) from the 1960s and earlier.

Like TV, it was intended that the camera feeds for CCTV surveillance systems were to be watched by a human operator. Realizing that operators needed to take breaks, sometimes missed something important that was portrayed on screen, or simply wanted to replay an event, they started recording the images. The first widely deployed platforms for this were VCRs, which continue to be used in many older installations and continue to be sold for entry security systems such as in single store retailers, gas stations, and the like.

In professional installations, VCRs have been largely replaced by disk-based systems called DVRs (Digital Video Recorders). Both VCRs and DVRs are similar to their home use consumer equivalents for video entertainment uses, recording and playing back television broadcasts.

DVRs are typically based on standard Wintel PC motherboard, disk drives, power supply and an ADC (analog-digital converter) or encoder, and running a version of Microsoft Windows.

Connectivity is via BNC connector ports on the back of the box to coax cabling for analog cameras. The ADC or encoder converts the analog video to digital format before it is processed by the DVR's CPU and recorded to internal disk. Many DVRs have one or more standard IP ports for connectivity choices. A small minority of DVRs leverage Linux or proprietary operating systems in place of Windows, but otherwise are configured and function in the same general manner. A very large industry manufactures DVRs, mostly from names unfamiliar to IT users.

A challenge that many security departments face today is related to staffing, resources, budget and new outside requirements for additional storage, better data quality, and the like. While those challenges parallel similar ones faced by IT, they differ in implementation for physical security.

We often have too many cameras (called *channels* when interfaced with DVRs) and not enough human operational staff in security departments, in part leading to the demand for new IP-based technologies to be developed and deployed because of multiple technology benefits. With these new systems, the surveillance application is transitioning from being TV-like to being a true video data application.

Foremost in the new digital world are IP cameras, which are managed by Video Management Software (VMS) applications. Many VMS applications actually record to disk first, prior to displaying on a monitor that human operators might watch. In fact, fewer operators actually watch live surveillance feeds today, as the "watching" is now done more and more by the cameras, by video analytic programs, or by human operators afterward once prompted to do so by the system.

The People

Traditionally, CCTV was largely installed and serviced by low voltage electricians at building construction sites, who followed schematics designed by Architect & Engineer (A&E) firms, and were specified in the plans for the building or facility. Low voltage electricians also installed access control systems, burglar alarms, fire detection, intercoms and public address systems.

Many of these low voltage electricians are now employed by dealer/installer organizations, and have grown IT capable and knowledgeable. However, this component of the physical security industry as a whole largely remains low-tech when it comes to sophisticated IT implementations involving setting up complex commodity server and storage configurations, VPNs, IP network troubleshooting, DNS/Active Directory services, subnet masking, OS tuning parameters, file system organization, and I/O tuning for specific application anomalies that are familiar to IT professionals.

Further, while physical security was largely entirely responsible for the video surveillance systems in an organization at one time, today this is changing. As more and more IP-based systems are deployed, a greater number of IT departments are being asked to participate and manage these systems in support of or in conjunction with physical security.

CCTV Technology

Many installed DVRs are considered FRUs (Field Replaceable Units) in IT terminology, meaning that if there is a failure, the entire DVR (video data and all) is typically shipped back to the manufacturer for repair and a spare DVR is used in its place. This is especially true of older DVRs, although newer versions may address this concern with field replaceable components.

Regardless of recent changes, many installed DVRs suffer from 3 common problems.

- *Storage capacity is usually limited to the sheet metal boundary of the DVR, and yet due to regulations and/or operational need, the customer is driven to retain more days of video (known as the retention period).*

- *Storage capacity also limits the resolution/frame rate (quality of recording), as increased resolution and frame rate demands reduce the retention capacity (how long the recordings can be kept) and vice versa. The trend to greater image recognition through the recording of high resolution megapixel images, combined with a need for more images per second (IPS), is a major driving force for additional capacity.*

- *Reliability issues caused by disk drive failures plague many installed DVRs because the I/O is continuous, and read/write conflicts result in excessive head travel. Often the disks used in older systems are not OEM quality, frequently not RAID protected (the IT industry standard form of data protection), and can cause over 50% of video surveillance system failures. Symptoms include missing video where the recording has stopped. Since the system continues to present images on a screen, the operator may not realize recording has stopped. Similarly, camera blackout - where the drive failure is catastrophic - can bring down the whole DVR and cause the screen monitors to go blank, is another common failure. Some newer DVRs address this, but not all do and many older DVRs remain in use that suffer from these challenges.*

These are the **3 Rs of Video for DVRs**: ***R****etention,* ***R****esolution/frame rate, and* ***R****eliability.*

Newer DVRs may allow field maintenance for failed drives, and may offer hot-swap components in some cases. They may also include RAID (Redundant Array of Independent Disk) data protection for improved reliability. However, RAID is still limited to the capacity of a single DVR, which results in a significant percentage of total capacity consumed by providing that support and thus adds significant cost and overhead.

Newer DVRs may also enable external recording capacity upgrades, most commonly by propriety, SCSI, USB or Firewire connections. As we will see when we discuss storage differences, this recording capacity continues to have limitations, not only restricted to a single DVR in most cases, but also in performance characteristics.

A remaining issue with DVRs is a fixed recording ceiling. While a typical DVR has 4, 8, 16, or 32 channels, the recording ceiling stays the same because as more channels are deployed, fewer frames per channel are recorded. This is due to the fact that the DVR receives analog signals, and needs to encode them into digital form in order to be able to write to disks. This encoding is a very CPU intensive process.

As a result, the quality of live video displayed on monitors is excellent as there is no processing needed for displaying analog signals. Unfortunately, this also means that the resolution/frame rate displayed on live monitors is often not what is actually recorded.

Further, DVRs are typically not managed in the IT network (via SNMP or other SRM monitors), and often provide no remote alerting to changes in camera/channel/stream flow. Again, the older or lower cost the DVR, the more likely these are to be the case.

TV Becomes a Video Data Application

In the beginning, recorded TV had limited perceived value since it was not searchable. It was also very cumbersome and time consuming to manually go back and look for something captured on tape. As a result, after some amount of calendar time had passed and the video was not referenced, the video tapes were usually recycled by being recorded over again and again. This resulted in the recorded video being destroyed as it was recorded over, and tapes gradually degrading with lower video quality over time.

The time period (known as the retention period) of how long tapes were kept before being recycled in this manner was set by the physical security organization according to a risk or threat assessment, or outside requirements. For casino gaming for instance, the requirements were often set by Gaming Boards, usually at 7 or 14 days. Other requirements were often set at 30 days retention. This made for a lot of tape management, and that eventually translated into a lot of disk storage for DVRs when they were transitioned in to replace the VCR platform.

However, it also helped condition many security practitioners to think of video as not being all that valuable, and declining further in value as it aged. Video was important if it caught a robbery in action, but 30 days of nothing much happening led to its perceived lack of importance. As a result, high availability systems, data backups and long term retention of video were not common like their counterparts are in IT. New security and non-security applications alike have begun to change this, each demanding longer retention periods and better video quality, but there remain vestiges of this attitude in many video surveillance system designs.

Video surveillance systems continue to grow exponentially in capacity, and not always for traditional security needs. For example, in some organizations this may be driven by Sarbanes/Oxley and HIPPA compliance requirements (after all, this is now *enterprise or corporate data*). Consider a retailer where a megapixel camera is able to record a credit card or

bank debit card number as it is used at a check stand. For some, that may become a potential SOX issue, since it is financial data that needs protection. While not a major impact on the market yet, challenges like this continue to push the need for increased retention and protection of video data.

In others, video data growth is often driven by risk mitigation. Legal counsel for retail stores are frequently requesting 2 years of retention, as the statute of limitations on *slip-and-fall* claims is 24 months in most U.S. states. We have also seen government, law enforcement and prison requirements grow well beyond this retention period for similar reasons, with corresponding increases in frame rate and resolution settings.

New uses of video data are emerging, for instance, in manufacturing and supply chain environments, where megapixel and/or HD IP cameras are increasingly installed over the production line. Production runs are recorded, and quality assurance is performed on the recorded video rather than spot checking the live run. This video data may be required to be kept for FDA and DEA audit purposes for pharmaceutical and medical equipment manufacturers, or by other agencies for food production and the like.

Real-time analysis is done by advanced cameras themselves, which are programmed with a couple of dozen different scenarios like *person fell down*, or *left behind object*, where someone has left a suitcase or other item. These are called *camera-based analytics*.

Historical analysis is done on the server side, using recorded video data for applications like *facial recognition* (e.g. *How many times have you seen this face in the last 6 months?*). This typically requires a growing database and matching video storage capacity.

While the price of storage per MB and TB continues to decline, the increase in demand for additional capacity continues to offset this, with larger video retention requirements growing. However, it is worth noting that the price of storage that physical security is willing to pay is significantly lower than that common in IT, especially since features like snapshots, replication, de-duplication and similar advanced capabilities are not perceived as being useful but are commonly built into IT systems.

The IP Platform

As an IT professional, if you are involved in a video surveillance project, it probably means that you are looking at a DVR (Digital Video Recorder) installation. It was also likely originally designed to work with analog cameras over coax cable, and installed and managed not by IT but by physical security. Or you may be looking at a more modern IP video surveillance solution that supports IP cameras over 1GbE IP Ethernet (or even POE). Although IT has largely left video surveillance up to the physical security department to buy and manage in the past, IT is becoming involved in more organizations today due to budget issues and the growing importance of recorded video.

For surveillance, IP systems come in several flavors:

1. *A Network Video Recorder (NVR), which is a heavily-tuned commodity server/storage combination, with proprietary video software, and private labeled by one of the DVR makers to be sold as a packaged solution;*
2. *A software-only NVR from one of the video management vendors, unbundled, to run on an existing customer infrastructure or on do-it-yourself server/storage platforms;*
3. *A video appliance, supporting a choice of video management system from a variety of vendors and Windows or Linux based operating system, integrated and ready to run out of the box with computer and storage integrated in a single, video optimized platform.*

As video encoding is done on the camera side, cameras have become more intelligent. NVRs and video appliances receive packets from IP cameras, process, and write them to storage subsystems. The processing can be very CPU intensive, as some amount of video decoding is required. Displaying video, which is usually done at the client side (a viewing station), is another very CPU intensive process.

Every DVR, NVR, and VMS has a unique I/O signature comprised of I/O size, I/O randomness, block allocation, chunking, and file system characteristics.

IP security cameras from a wide array of vendors cover an even wider array of capabilities, including piezoelectric motor driven, PTZ (pan-tilt-zoom) functionality, FLIR (Forward Looking Infrared or thermal, to see in the dark), outdoor ruggedized, dome (fixed), motion activated, etc. Cameras at the receiving end are often referred to as *channels* (in DVRs) or *streams* (often in NVR and video appliance terms) in physical security terminology.

Depending on camera choice, video streams consist of images in a compression CODEC (Coder/Decoder) like MPEG, MJPEG, MPEG4, JPEG 2000 or H.264.

Further, cameras and CODECs determine resolution, which is measured as CIF, 2CIF, 4CIF, Megapixel, 2Megapixel, 3Megapixel, and so on up to 10 Megapixel today. Take note that what is recorded can sometimes be better resolution than what is seen on a live monitor.

Bit-rates from each channel will vary depending on the CODEC, resolution, motion-activation, and frame rate settings, which are typically chosen at setup time.

Other Important Aspects of the Video Surveillance Application

The typical video surveillance environment is 7x24x365 constant streaming with little or none of the overnight batch operations common in IT. There are some exceptions depending on the operational circumstances for specific application environments. For example, retail stores can have a quieter night time phase, while ATMs and casinos typically do not.

A video surveillance system has a single opportunity to capture video frames. Otherwise, the imagery is lost forever, as there is no re-transmission opportunity in this always-on, live recording application. Fault resilient configurations cover almost all situations, but occasionally we see a high availability (HA or non-stop) data protection requirement for the most demanding environments. We may see this HA need in highly sensitive embassy operations, gaming, and often at key facilities like a customs hall in an airport. This is where dual data-path architectures are deployed, otherwise uncommon in physical security.

Video surveillance streaming requires massive bandwidth between the camera and the recording platform over Ethernet, either on a dedicated network or via VPN. Hence, these tend to be over LANs and not WANs for deployment. We also commonly see bit rates exceeding 10Mbps for a single, standard IP camera, and for a megapixel camera a single stream can be over 30Mbps. This has significant impact on network infrastructure and design, and limits remote storing of video for real-time monitoring applications.

Once the video stream is set up, it is all about I/O operations. There are massive disk throughput requirements from the CPU to the storage media that operate non-stop. Disk latency restrictions and network bandwidth usually mean that storage is co-located with the compute resource, and idle, powered down disk arrays for energy conservation are not commonly used in this application.

Most video management software uses a file system in the host OS (mostly NTFS based), and hence ultimately translates into SCSI (DAS) or iSCSI (SAN) block-based storage requirements. We commonly see I/O rates of up to 30 IOPS per camera. Spindle count and I/O parallelization are key factors of disk throughput performance, so storage scalability is critically important.

File-based Network Attached Storage (NAS) is almost never used as a primary recording method in video surveillance, since it adds an unnecessary management layer to complicate and slow performance, but can be used effectively as a 2nd tier medium. NAS can also be useful for archiving short clips of video usually kept as evidence, often referred to as *clips of interest*. Protecting those clips from accidental or intentional editing or deletion is a key requirement for the video storage system.

Fibre Channel (FC) based storage is also very rare in this environment, and can be indicative of the intensive I/O nature of the application. When the installation is supporting many cameras, users have sometimes turned to FC thinking that this was their only choice to keep up with the I/O streams. However, FC is cost justifiable only in rare cases since physical security budgets typically do not support the price per TB commonly demanded for IT grade, FC SAN storage. Similarly, while a single project may required hundreds of TBs of capacity, such as in a housing authority surveillance system or for a public transit system, the video is typically stored locally at multiple, smaller installations rather than in a central repository for real-time access. A central repository may be used for archiving video, storage of clips of interest, and the like, but is not typically the primary storage used.

Storage capacity requirements of video surveillance data can be a significant challenge, far exceeding typical web servers and databases. A small commercial installation can be as large as 10TBs, and in midrange systems 100TBs is common. Large systems of 1– 4PB (Petabytes) can usually be found in casino gaming, municipal surveillance, and some government installations, but again may be distributed in several or multiple sites due to bandwidth issues.

There is no *cumulative history* in this application, as there would be in building a financial record in an IT system. In video surveillance, the recording media is continually recycled in a loop, varying only in length of retention. This causes excessive fragmentation and head travel if one is not careful in system design. Systems typically do not allow downtime for management processes like defragmentation, and most surveillance systems must run unattended without operator tuning and system maintenance. In combination with the recycling of storage, this is the reason the application supports an *evidence archive* feature, where *clips of interest* as mentioned earlier are copied to another location (usually over a network file protocol like NFS or CIFS) in a non-real time manner.

There is no *backup, snapshot* or *data de-duplication* function used in this application. The volume of video data is simply too large and time consuming to make a backup in a traditional sense in IT, and video software capabilities lack the finesse and horsepower required for de-duplication techniques that are growing popular in IT data file environments.

Sometimes users will deploy a multi-cast camera approach where they create two streams and make two real-time copies in separate locations. These copies differ in resolution and frame rate because they are traveling over different distances where cost/capability restrict the bandwidth available. This is being done for *last resort* or disaster purposes, and is not common today.

The purpose of multi-cast is not to reconstitute the entire video data set, but to capture the last moments of an event, even if the video is at a much reduced resolution and frame rate. An example would be a bank robbery where the bandits stole the local DVR on their way out the door, a not unheard of occurrence by savvy thieves. If a second stream was created offsite, there would be at least some evidence and a video record of the incident.

Configuration changes of a video surveillance system vary, as with a database application. However, significant changes happen every time there is a physical modification to the camera count, camera settings, or when the retention period is extended. Often a government regulation or corporate policy change is mandated, increasing the number of days the video data needs to be kept or increasing the required resolution and frame rate, and having a considerable impact on the available storage but not the systems' performance for camera support.

The Data Model

At first glance, one might think that video surveillance means cameras, and that cameras mean sequential streams of video to be recorded. This is true for a single camera, but not for more than that.

Though it is counterintuitive, it turns out that video surveillance data is actually random. This is because we are taking frames from multiple cameras simultaneously, and each camera has a variable bit rate caused by motion activation, key frames, scene changes or other factors and settings. Further, this is only the first level of randomization. Another level is introduced by the placement of these video streams into a file system (NTFS for Windows systems, or LFS/EXT for Linux systems).

Therefore, by the time the I/O drops out of the operating system, the net result is mostly large, random I/Os in block form. There are also often small random I/Os of metadata mixed into the stream. Hence, the I/O request size can range from a few KBs to hundreds of KB each. Most of the metadata is small in size, while the video data I/Os have a big range depending on the VMS applications themselves, which can vary widely.

Continual loop recording means constant recycling of the same set of blocks on disk. The result is often excessive disk fragmentation, occurring when the write algorithm is either *write anywhere* or *disk head/seek time optimized*. Further, it means that the sequential nature of data before it is committed to rotating media and block allocation strategies are critical factors.

The lifecycle of video surveillance data also demands careful consideration on how this application potentially impacts deployment on existing IT infrastructure. There is often a need for separate networks or isolated dedicated storage in order to avoid dramatically impacting database application performance with the VMS application.

"The Achilles heel of any modern storage system is disk performance when the application has a randomly-accessed working set larger than cache memory" -- SNIA 2007.

1. *Structured data like databases consist of small I/Os, random, mostly reads, perhaps 60%/40% writes and reads.*

2. Unstructured data like BLOBs (External Storing of Binary Large Objects) consist of large I/Os, sequential, lopsided to either almost all writes or reads.

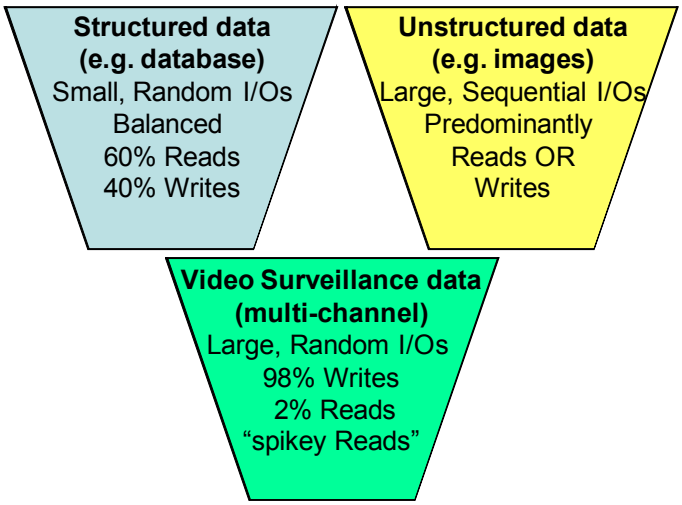

Video Surveillance Data Type

Video surveillance data is neither *small/random* nor *large/sequential*. Instead, it is a *large/random* I/O workload.

Consider that video surveillance streams are typically 98% writes and 2% reads. However, the reads tend to happen all at once, so they are very “spiky”. For example, when there is an incident and everyone jumps onto the system at the same time to view (read) video data, we then see such a spike of reads from recorded video while recording from cameras continues unabated.

These are the times of peak performance loads. While the spiky reads are occurring, the system must continue capturing streaming writes and must maintain zero frame loss. Here, constant low latency is critical, unlike common or traditional IT applications which often have some leeway.

Loop recording creates *fragmentation and randomization,* which in turn causes poor performance over the long operational periods of video surveillance. Since video surveillance systems are typically 24x7x365 operations and do not have large, skilled IT support staffs and skill sets, the system must be designed to minimize fragmentation, randomization and any other results that would require operator intervention. This is not common in IT system design, where it is assumed that offline maintenance will routinely occur.

Intransa’s unique Video Data Management and Retention™ (VDMR) technology includes patented Video Surveillance Optimization (VSOP). VSOP mitigates the performance problems by re-sequencing random I/Os before they are laid down on the media. This ensures a more efficient write function and significantly better disk utilization.

However, the real benefit is that it also makes for a more efficient read operation, so that whenever reads occur, they don’t impact the streaming writes. In this way Intransa video

storage solutions are able to avoid frame loss that plagues traditional IT storage solutions used for video surveillance deployments.

Purpose-Built Intransa VideoAppliance™

Tier 1 performance (VSOP) with Tier 2 SATA for cost and scalability.

1. Use cases include IP cameras for new installations and use with existing DVRs to extend the life of CCTV.
    a. *Hosting a VMS to support streaming IP cameras (I/O engine coupled with Compute)*
    b. *Solving the 3Rs of Video with DVR streaming (iSCSI)*
    c. *Support for DVR clips/evidence archiving (NFS, CIFS)*

Intransa VideoAppliance™ simultaneously provides iSCSI, NFS, CIFS, and is a tightly coupled SAN/Server. Each supports Video Management Software (VMS) leveraging patented Video Data Management & Retention™ technology. Rack mount and pedestal configurations are available.

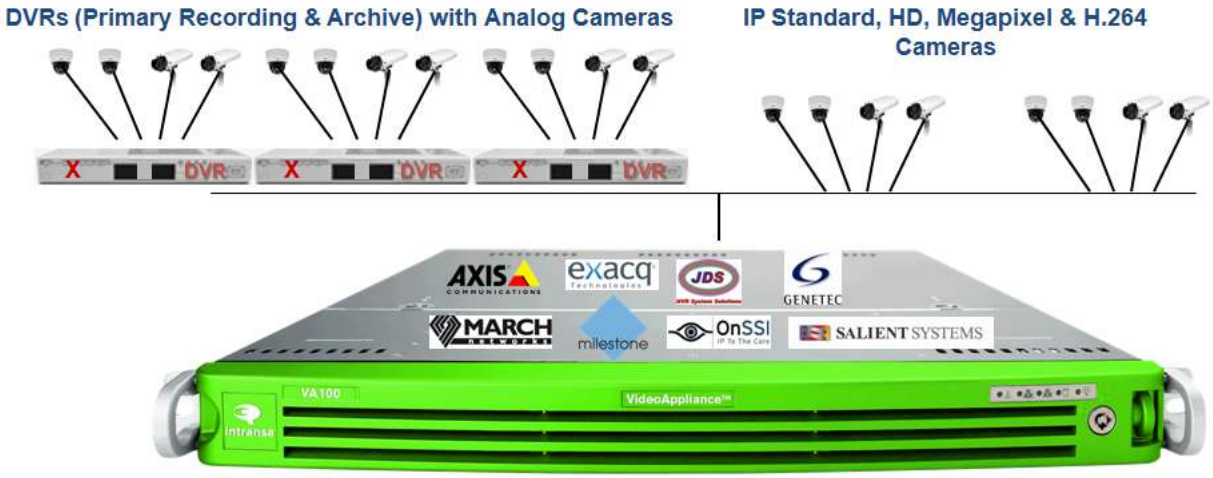

2. Ease of implementation:
    a. *System configuration is driven by camera compression CODEC, resolution, frame rate, amount of motion activation, and retention period*
    b. *Setup of the application is camera-specification driven. Calculators are an industry standard way of identifying setup requirements during installation*

Intransa VideoAppliance™ VDMR technology includes the Video System Administrator (VSA) GUI that provides automatic configuration and provisioning based on camera specifications.

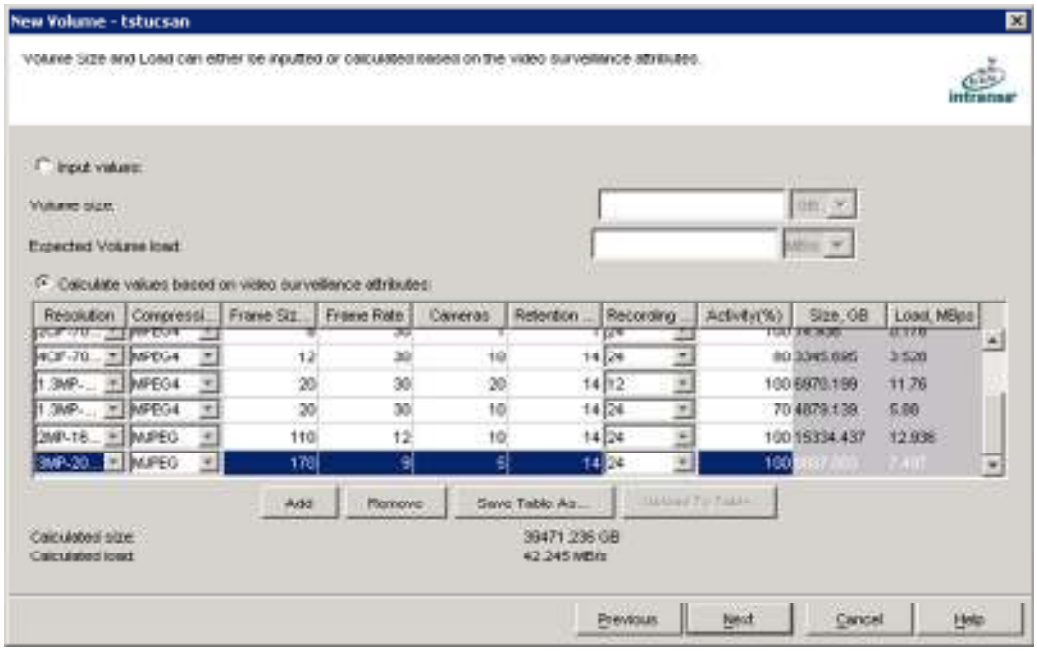

Intransa VideoAppliance™ provides IT staff full access via a CLI and GUI, such as the VSA screen used in VA100st and VA200st storage appliances

3. Manageability and Unattended Operation
    a. *Threshold setting and detection of camera streams at I/O level*
    b. *Alerting via SNMP (full MIB for use with Enterprise Management consoles like IBM Tivoli® and HP OpenView™)*
    c. *Email notification*

Intransa VideoAppliance™ provides these capabilities in the standard package.

4. Server Settings
   Commodity servers used in do-it-yourself configurations are tuned for IT applications, and must be re-tuned as I/O engines for effective use in video surveillance applications.

   Such servers can be either Windows or Linux based, but most video servers host Windows-based applications. For example, Windows registry settings are adjusted for best practices and integration parameters are set for plug and play, resulting in a system with known performance characteristics.

   Intransa VideoAppliance™ is preconfigured and pre-tuned for VMS support.

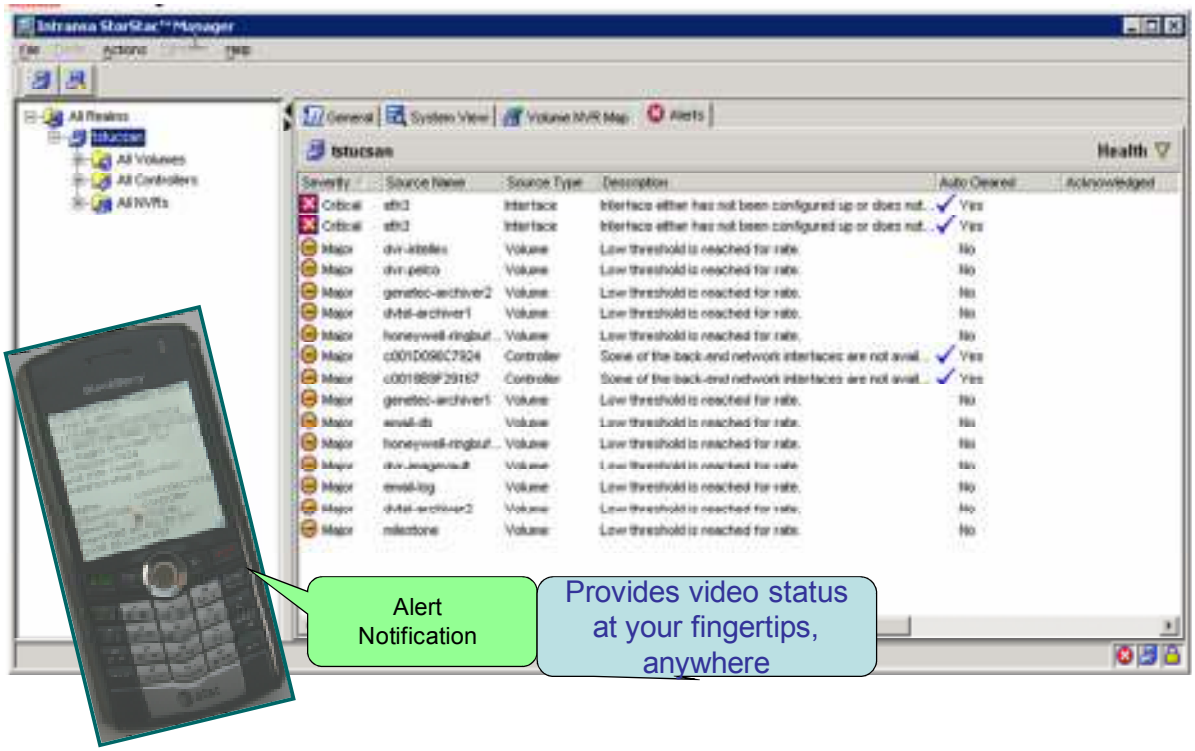

5. Security of the System
    a. *Must be managed like any other Windows box in corporate network*
    b. *Must be updatable as required with Windows Service Packs, patches, and Security Updates*
    c. *Must be able to run corporate standard firewall and virus protection as required*

Intransa VideoAppliance™ fits into the Corporate Security schema like any other Windows box on the corporate network, and is Microsoft® WHQL Certified (re-certification refresh in process).

6. Risk Mitigation through Interoperability
    a. *Intransa VideoAppliance™ Certified Program*
    b. *Interoperability / integration with Video Management Software providers*
    c. *Software*
    d. *Cameras*

Intransa VideoAppliance™ has been certified with several hundred security products, assuring risk-free interoperability and reducing implementation and support challenges. Multiple VMS choices are also available, with preloaded media kits in each appliance.

7. System Extensibility
    a. *Capacity expansion via advanced IP-SAN-in-a-box technology*
    b. *VMS interfaces to a LUN, and LUNs can be cascaded*

Intransa VideoAppliance™ comes with a 1U or pedestal 4 disk drive bay base unit and can be expanded with one or more additional 12 drive bay 2U units, and that in turn can cascade to additional block level capacity in other Intransa VideoAppliances.

8. Tuned for Video Surveillance I/O
    a. *Must be tuned for the unusual large, random I/O profile of video surveillance*

Intransa VideoAppliance™ has Windows tuned for throughput performance and disk I/O latency, and contains patented Video Data Management & Retention™ technology that includes VSOP (Video Surveillance OPtimization) to improve streaming performance, and re-sequences data streams for efficient organization on disk media.

9. Support Services for Video Surveillance Systems
    a. *Hardware and software must function as a cohesive application*
    b. *Cooperative arrangement with certified Video Management Software providers is provided*
    c. *VMS Software in our Support lab*
    d. *Level 2 support*

Intransa has a follow-the-sun support organization that provides 24x7x365 support, capable of 4 hour on site response worldwide.

Conclusions

Physical security overall and video surveillance in particular differ greatly from traditional IT environments. Many features common in IT are not used, while other important ones are unique to video surveillance. Yet for all of the differences, video surveillance is a growing requirement for many IT departments to understand, engage and support.

Intransa VideoAppliance™ is an ideal platform to bridge the technological and organizational differences between IT and physical security.

The complete Intransa VideoAppliance™ family meets a wide range of needs for physical security. It's why we were named *"2009 New Product of the Year*" at ASIS International in Anaheim, California. And it's why you should remember Intransa, the VideoAppliance Company®, for your next physical security project.

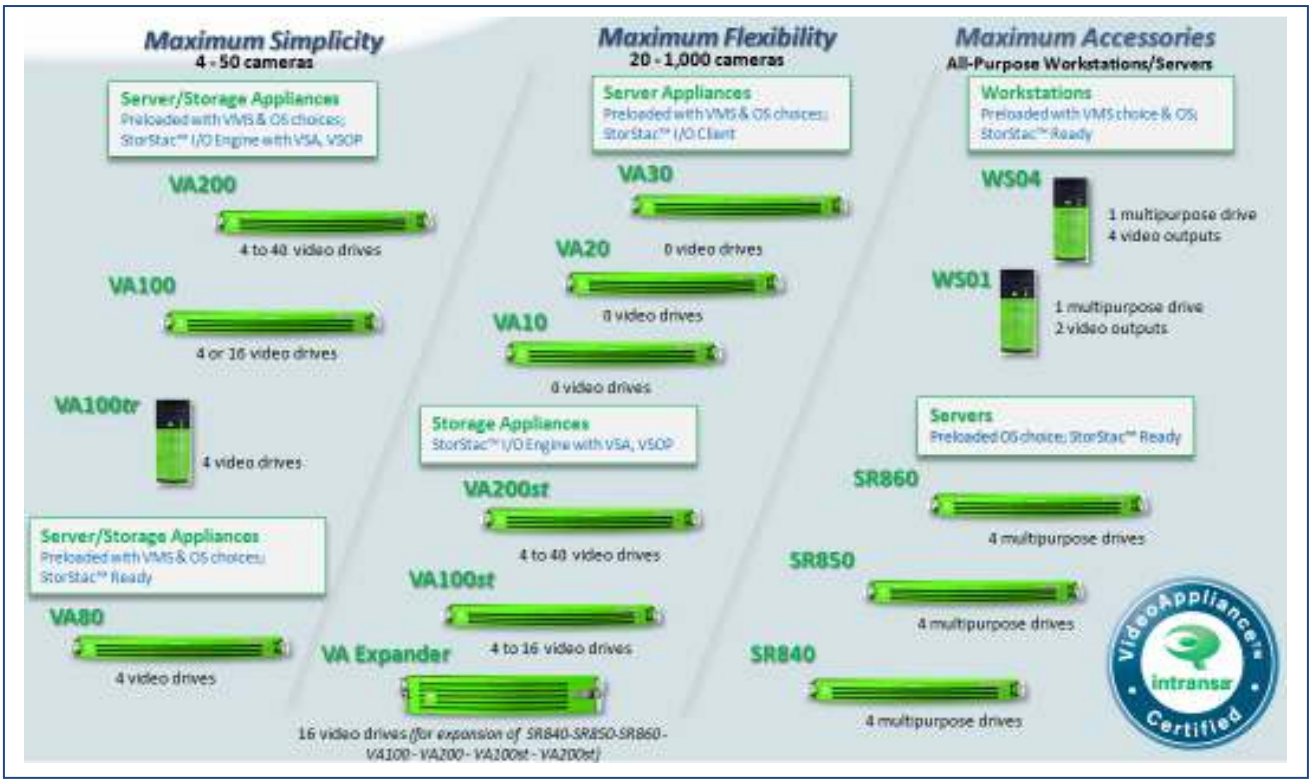

To learn more, please contact Intransa or an authorized dealer or integrator today.

Intransa, Inc. / www.intransa.com / www.videoappliance.com
10710 N. Tantau Avenue, Cupertino, CA 95014 USA

Toll free 866.446.8726 **International** 408.678.8600 **Email** sales@intransa.com