



The Curious Case of Video Surveillance

White Paper

V1.1





Table of Contents

Introduction	3
Background	3
The People	3
CCTV Technology	4
TV Becomes a Video Data Application	5
The IP Platform	6
Other Important Aspects of the Video Surveillance Application	7
The Data Model	9
Purpose-Built Intransa VideoAppliance™	11

INTRODUCTION

Video Surveillance and Security are unlike applications typically found in IT environments. Video surveillance systems are real-time, resource intensive systems that run 7x24x365 in a streaming mode, where the processing power, network bandwidth, disk throughput, and storage capacity requirements are extreme. A proper design of servers, networks, and storage is the key to successful video surveillance systems.

Video Surveillance is often like a 911 system. A bad day in Security is when someone is injured or killed, and is especially bad if the video system capturing the event has failed.

Architecting a reliable surveillance system without excessive cost 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 storage requirements.

BACKGROUND

This application originated from analog CCTV (Closed Circuit TV).

Like TV, it was intended that the camera feeds were to be “watched” by a human. Realizing that humans needed to take breaks, sometimes missed something, or simply wanted to replay an event, they started recording the images on VCRs, which were later largely replaced by disk based systems called DVRs (Digital Video Recorders). These are mostly embedded Windows PCs with BNC connector ports on the back of the box, and “encoders” to convert analog to digital inside the box before it is processed by the CPU and recorded to disk.

Today, we often have too many cameras (called “channels”) and not enough humans, leading to new IP-based technologies being developed and deployed. With these new systems, this application is transitioning from being TV-like to being a Video Data application.

Many Video Management Software (VMS) packages actually record on disk first, prior to displaying on a monitor that humans might watch. In fact, fewer humans actually watch live surveillance feeds, as now the watching is done by the cameras, by analytic programs, or by humans afterward once prompted to do so.

THE PEOPLE

Traditionally, CCTV was installed and serviced by the “low voltage electricians” at building construction sites, who followed schematics designed by Architect & Engineer

(A&E) firms, and specified in the plans for the building. Low voltage electricians also installed access control systems (“card key access”), burglar alarms, fire detection, intercoms and public address systems.

Many of these dealer/installers have now become IT capable. However the industry as a whole largely remains low-tech when it comes to sophisticated IT implementations involving setting up VPNs, network troubleshooting, DNS/Active Directory, subnet masking, OS tuning parameters, file system organization, and I/O tuning for specific application anomalies.

While Security largely was 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 or conjunction with Security.

CCTV TECHNOLOGY

Many installed DVRs are FRUs (Field Replaceable Units), meaning that if there is a failure, the entire DVR (video data and all) is typically shipped back to the Security 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.

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

- Storage capacity is 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 (**Retention** Period).
- Storage capacity also limits the **Resolution**/frame rate (quality of recording), as more resolution and frame rate reduce the Retention (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, nobody realizes recording has stopped. Similarly, “camera blackout” where the drive failure is catastrophic, bringing down the whole DVR and causing the screen monitors go blank, is another common failure.

These are the “**3 Rs of 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.

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

A remaining issue with DVRs is a fixed recording ceiling. While a typical DVR has 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. This also means that the resolution/frame rate displayed on live monitors is often not what is recorded.

Further, DVRs are typically not managed in the network (via SNMP or other SRM monitors), and often provide no remote Alerting to changes in camera/channel/stream flow. Again, the older 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, and it was very cumbersome and time consuming to manually go back and look for something captured. As a result, after some amount of calendar time had passed and the video was not referenced, the video tapes were usually recycled. This resulted in the recorded video being destroyed as it was recorded over.

The time period (Retention Period) of how long tapes were kept before being recycled was set by the organization according to a risk or threat assessment. For casino gaming for instance, it was set by Gaming Boards usually at 7 or 14 days, or whereby Security organizations commonly set retention at 30 days. This made for a lot of tape management, and that eventually translated into a lot of disk storage for DVRs.

Video Surveillance systems continue to grow exponentially in size, and not always for traditional security needs. For example, in some organization it 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 a couple of years in most states. We have also seen government and prison requirements grow well beyond this retention period for similar reasons.

New uses of video data are emerging, for instance, in manufacturing where HD cameras are installed over the pharmaceutical 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 is kept for FDA and DEA audit purposes like searching, counting, and other analysis.

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 object”, where someone has left a suitcase or something, etc. These are called camera-based analytics.

Historical analysis is done on the server side, using recorded video data for applications like Facial Recognition (how many times have you seen this face in the last 6 months?). This typically requires a growing database.

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.

THE IP PLATFORM

As an IT person, if you are involved in the project, it probably means that you are looking at a DVR (Digital Video Recorder) installation which was originally designed to work with analog cameras over coax cable and installed and managed by Security. Or you may be looking at an IP video surveillance solution that supports IP cameras over 1GbE Ethernet (or POE). Although IT has largely left video surveillance up to the security department to buy and manage in the past, IT is becoming involved in more organizations today due to budget issues and other factors.

For surveillance, “IP systems” come in several flavors:

1. A Network Video Recorder “NVR”, which is a heavily-tuned commodity server/storage, with proprietary 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 NVR makers product, unbundled, to run on an existing customer infrastructure;

3. A 3rd party Video Management Software (VMS) package (Windows or Linux), running on an Intransa VideoAppliance™ or a do-it-yourself server/storage platform.

As video encoding is done on the camera side, cameras become more intelligent. NVRs 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 (“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 (point-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 or streams.

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

Further, cameras and CODEC determine Resolution, which is measured as CIF (sometimes written SIF), 2CIF, 4CIF, Megapixel, 2Megapixel, 3Megapixel, and so on up to 10 Megapixel. Take note that what is recorded can sometimes be better resolution than what is seen on live monitor.


Bit-rates from each channel will vary depending on CODEC, Resolution, motion-activation, and Frame rates, which are chosen at setup time.

OTHER IMPORTANT ASPECTS OF THE VIDEO SURVEILLANCE APPLICATION

The typical video surveillance environment is 7x24 constant streaming with little or none of the “overnight batch” operations common in IT. There are some exceptions depending on the operational circumstances. For example, retail stores can have “night time” but ATMs and casinos cannot.

The system has 1 (one) opportunity to capture video frames. Otherwise, they are lost forever, as there is no re-transmission opportunity in this application. Fault Resilient configurations cover almost all situations, but occasionally we see a high availability (HA or non-stop) requirement. We may see this 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.

Video surveillance streaming requires massive bandwidth between the camera and the recording platform over Ethernet, either on a dedicated network or VPN. Hence, these



tend to be over LANs and not WANs. We commonly see bit rates exceeding 10Mbps for single, standard IP camera, but for a megapixel camera single stream can be over 30Mbps.

Once the stream is set up, it is all about I/O. There are massive disk throughput requirements from the CPU to the storage media. Disk latency restrictions usually mean that storage is co-located with the compute resource.

Most video management software uses a file system in the host OS (mostly NTFS), and hence ultimately translates into SCSI (DAS) or iSCSI (SAN) block-based storage. 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 important.

File-based Network Attached Storage (NAS) is almost never used as a primary recording method, 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.

Fibre Channel (FC) based storage is rare and can be indicative of the intensive I/O nature of the application. When the installation is supporting many cameras, users have 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.

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.

There is no “cumulative history” in this application, as there is would be in building a financial record in an IT system. In video surveillance, the recording media is continually recycled in a loop. This causes excessive fragmentation and head travel if one is not careful. This recycling is the reason the application supports an “evidence archive” feature, where “clips of interest” are copied to another location (usually over a network file protocol like NFS or CIFS).

There is no “backup” or “snapshot” 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. Sometimes users deploy a multi-cast camera approach where they create 2 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. This 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. If a 2nd 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, but changes happen every time there is a physical change 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, having a considerable impact on the storage but not the performance for camera support.

THE DATA MODEL

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

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. Further, this is only the 1st level of randomization. The 2nd level of randomization is introduced by the placement of these streams into a file system (NTFS for Windows systems, or LFS/EXT for Linux systems).

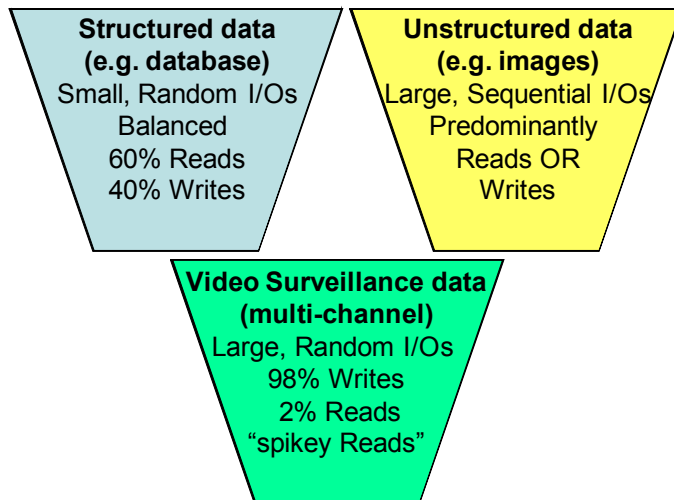
So, by the time the I/O drops out of the Operating System, the net result is mostly Large, Random I/Os in block form. There are also some Small Random I/O of metadata also mixed into the stream. Hence, the I/O request size can range from a few KBs to hundreds of KB. Most of the metadata are small in size, while the video data I/Os have a big range depending on the VMS applications themselves.

Continual loop recording means constant recycling of the same set of blocks on disk. Excessive fragmentation occurs if the write algorithm is “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 could impact existing IT infrastructure. Perhaps there is a need for separate networks or isolated dedicated storage to avoid dramatically impacting a 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 database consist of small I/Os, random, mostly Reads, perhaps 60/40 to Writes.
2. Unstructured data like BLOBs 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, it is Large/Random I/O.

Consider that streams are typically 98% Writes, 2% Reads. However, the 2% Reads tend to happen all at once, so the Reads are very “spikey”. 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.

These are the time of peak performance loads. While the spikey Reads are occurring, the system must continue capturing streaming Writes and must maintain zero frame loss. Here, constant low latency is critical unlike other traditional IT applications.

Loop recording creates Fragmentation and Randomization, which in turn, causes poor performance.

Intransa’s unique Video Data Management and Retention technology includes Video Surveillance Optimization (VSOP), patent-pending 2008. 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 better disk utilization.

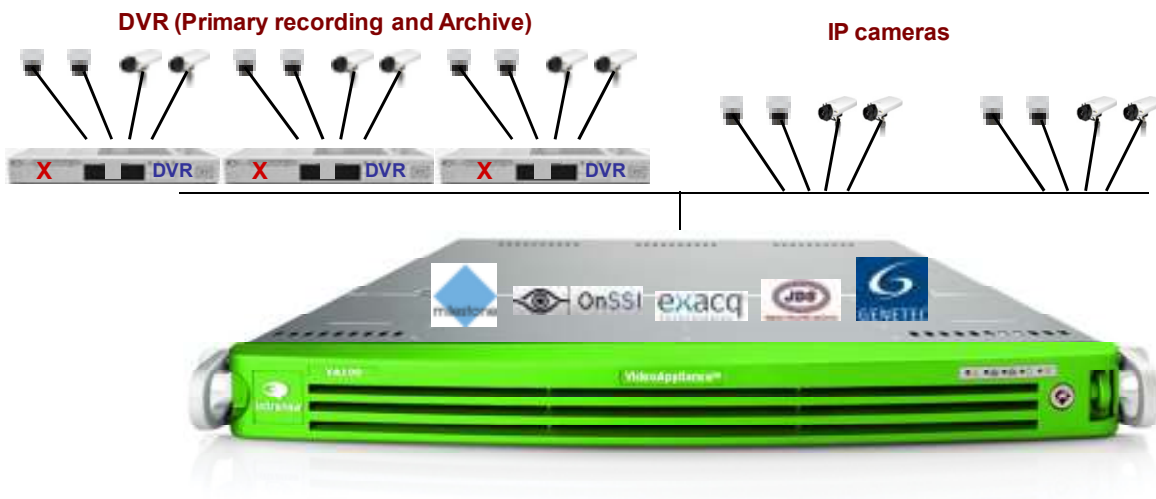
However, the real benefit is that it also makes for a more efficient Read operation, so that whenever Reads occur, they don’t impact the streaming Writes, and in this way are able to avoid frame loss.

PURPOSE BUILT Intransa VideoAppliance™

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

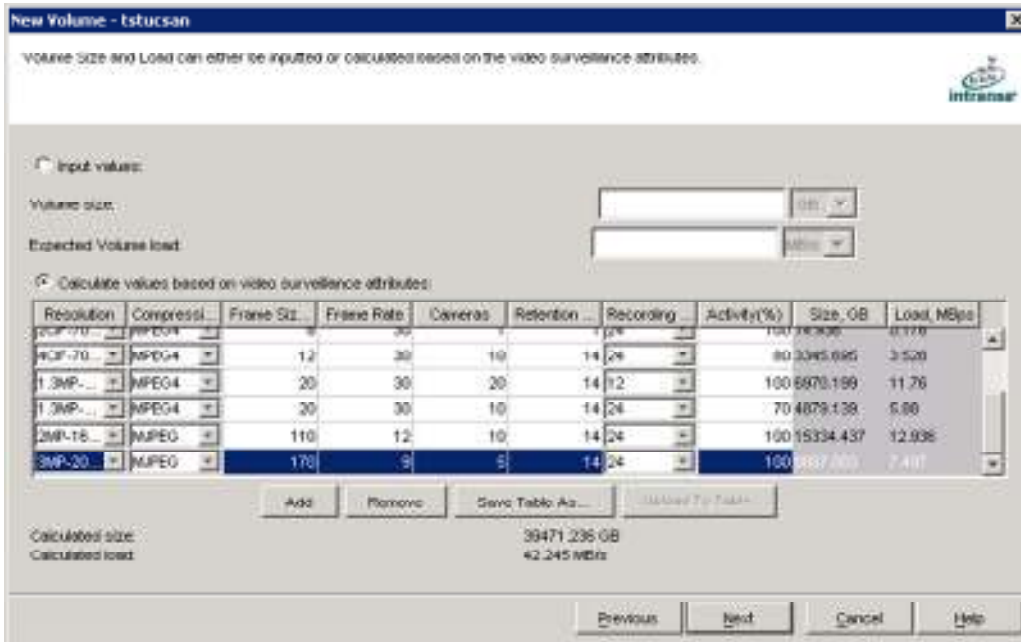
1. Use cases include IP Cameras and use with existing DVRs.
 - a. Hosting VMS to support streaming IP cameras (I/O engine coupled with Compute)
 - b. Solving the 3Rs of DVRs with DVR streaming (iSCSI)
 - c. Supporting DVR clips/evidence Archiving (NFS, CIFS)

Intransa VideoAppliance™ simultaneously provides iSCSI, NFS, CIFS, and is a tightly coupled SAN/Server supporting Video Management Software (VMS).



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™ contains the Video System Administrator (VSA) GUI that provides automatic configuration and provisioning based on camera specifications.

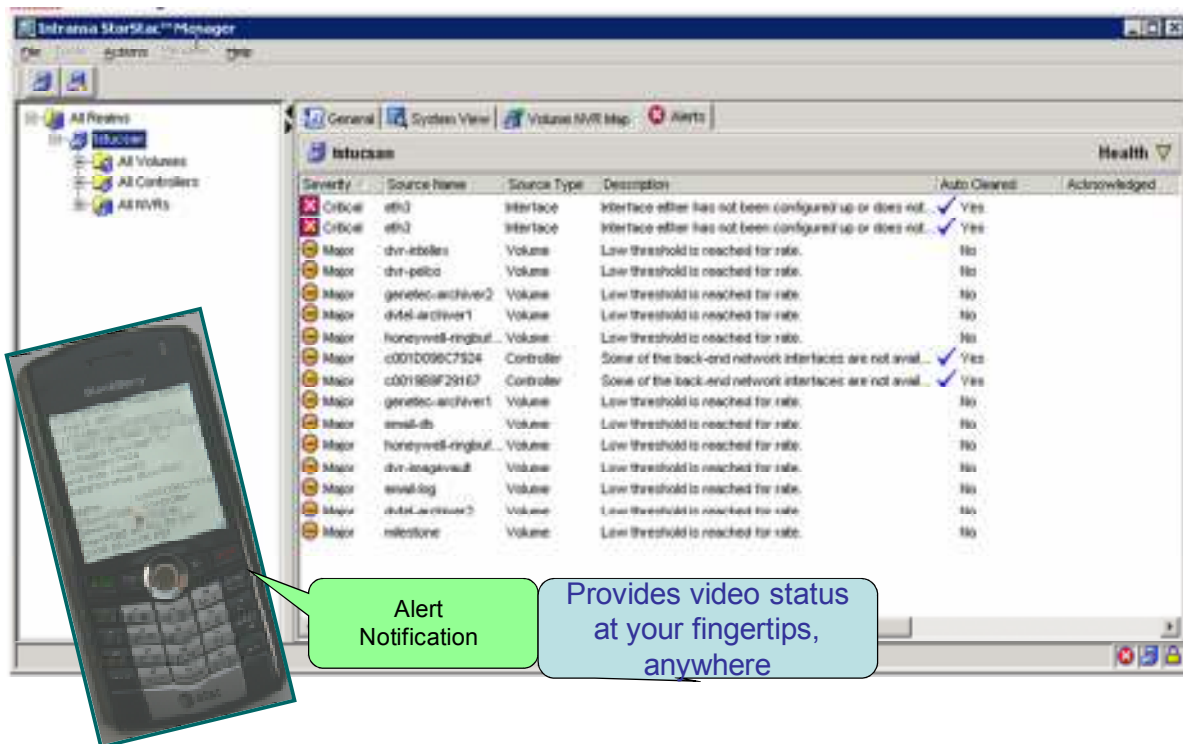


Note: Intransa VideoAppliance™ provides IT staff full access via a CLI and GUI

3. Manageability

- 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 are tuned for IT applications, and must be re-tuned as I/O engines for use in video surveillance applications. Servers can be either Windows or Linux, 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 Security 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 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



Intrinsa VideoAppliance™ has been certified with several hundred security products, assuring risk-free interoperability and reducing implementation challenge.

7. System Extensibility

- a. Capacity expansion via “SAN in a box”
- b. VMS interfaces to LUN and LUNs can be cascaded

Intrinsa VideoAppliance™ comes with a 4 disk drive bay base unit and can be expanded with an additional 12 drive bay unit (total of 16 drives), and that in turn can cascade to additional block level capacity in other Intrinsa VideoAppliances.

8. Tuned for Video Surveillance I/O

- a. Must be tuned for the unusual Large, Random I/O Profile of Video Surveillance

Intrinsa VideoAppliance™ has Windows tuned for throughput performance and disk I/O latency, and contains Patent-pending VSOP (Video Surveillance OPTimization) technology that improves 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 must be provided
- c. VMS Software in our Support Lab
- d. Level 2 support

Intrinsa has a follow-the-sun Support organization that provides 24/x/7/365 support, capable of 4 hour on site response worldwide.



Intransa, Inc.

Corporate Headquarters – 2870 Zanker Road, MS 200, San Jose, CA 95134 USA
408.678.8600 | 866.446.8726 | www.intransa.com | www.videoappliance.com | sales@intransa.com

© 2009, Intransa, Inc. All rights reserved. Intransa, the Intransa Logo, Intransa VideoAppliance, the Intransa VideoAppliance Certified logo, VDMR, and VSOP are business names and trademarks of Intransa. All others are the property of their respective holders and owners.