HITACHI
Inspire the Next

WHITE PAPER

The Evolution and Benefits of Hitachi Visualization Platform and Its Video Management Platform for Today’s Video Surveillance Architectures

By Hitachi Data Systems

July 2015

Contents

Executive Summary

This white paper takes a look at the evolution of video management platforms in the industry, presenting video management platforms in a historical context along with a detailed technical discussion and benefits of the approach. Video management platforms were created to address a fundamental need at the high-end of the market for a highly available, robust, converged hardware stack that is optimized for video processing for large numbers of cameras. The overall design is such that cost and performance are balanced with primary considerations for high throughput, redundancy and a fault-tolerant architecture that is unmatched in the industry.

To support public safety solutions that foster safer, smarter, more efficient communities through connected intelligence, Hitachi Data Systems has launched Hitachi Visualization. Hitachi Visualization helps law enforcement and emergency management personnel prevent and collaboratively address public safety situations. The solutions do this by integrating the capture, management and analysis of real-time information. They are part of the broad commitment by Hitachi Data Systems to Social Innovation in order to make society safer, sustainable, healthier and more comfortable by leveraging the Internet of Things, big data analytics, and IT infrastructure technologies.

They consist of hardware solutions known as **Hitachi Visualization Platform (HVP)** and an integrated map-based software platform called **Hitachi Visualization Suite (HVS)**, which enables correlation of disparate data and video systems geospatially. They are the flexible, searchable and visual solutions that government agencies, law enforcement and emergency management have been waiting for to enable campus and citywide data and video analytics and surveillance with public and private entity integration.

Hitachi Visualization Platform video appliance solutions, edge gateway devices, and specialized camera housing units allow organizations to intelligently collect, share and analyze information with built-in wireless networked data and video feeds.

This paper also examines the capabilities of the HVP video appliance solutions, HVP video management platforms: HVP 150 video management platform, HVP 500 video management platform, and HVP 1000 video management platform.

Introduction

As you're walking into a Las Vegas casino you automatically spot hundreds of black domes scattered all over the entire ceiling. It is no surprise that casinos have hundreds, if not thousands of cameras overlooking every inch of their real estate. They monitor every gaming table, slot machine, hallway and room. These cameras record 24/7/365 and are vital to maintain overall security and operational awareness of the business.

In some environments, like airports and government buildings, there are mandates that drive retention periods, resolution, and frame rate. In the case of a casino, it is strictly a business need. If a human eye can only detect somewhere around 15 frames per second (FPS), then our security system must be better than a human eye. Casino security must compensate for human deficiencies and have a system in place that would allow them to catch quick hand movements or any signs of double-dealing. Otherwise, it could result to significant monetary losses.

Typically, these casinos utilize megapixel cameras and record video at 30 FPS. H.264 compression is a common encoding mechanism to lower the network bandwidth and storage requirements. It's possible that cameras looking at mostly static images like doorways will only record on motion. It is also possible that some cameras support analytics to monitor directional flow of traffic or have facial recognition support to monitor casino's black list players.

Some key personnel have access to every single camera on the casino floor. Other, lower ranking personnel may only look at specific areas and may not be able to pan, tilt, or zoom the cameras that have that capability.

None of these aspects of video surveillance mechanisms would be easy to maintain without a digital video management system (DVMS) or VMS for short. **A VMS is camera management system software that allows centralized control of video cameras on the network.** It is essentially the brain of the surveillance environment. A VMS allows administrators to set frame rate, retention time, access, pan/tilt/zoom capabilities, analytics, access control and user rights of all cameras on the network.

VMS is a software suite, which means it requires servers for processing and storage for video retention. A typical video management system is mostly composed of a *management* server and *archiver* servers. A management server is the brain of the VMS. It controls every aspect of the software suite, including access rights, camera settings, storage parameters, and so forth. The archivers act as digital video recorders (DVRs), which record all video from the cameras that are attached to them.

Figure 1. Typical VMS Architecture

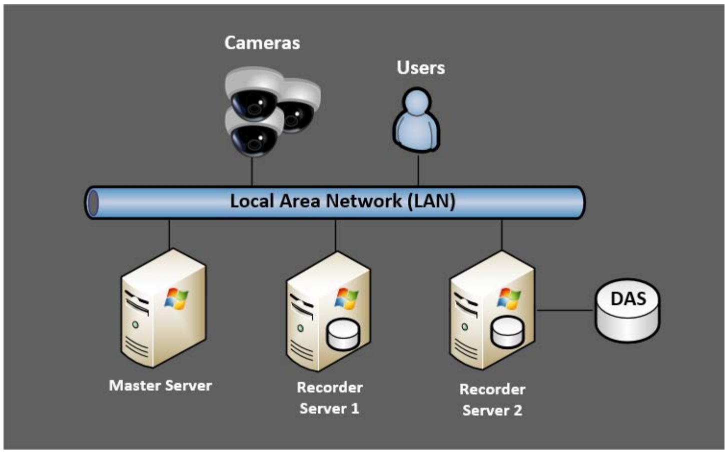

Looking at Figure 1, we see a typical VMS installation environment. Based on this figure, we see cameras, users, master server and recorders on the same local area network (LAN). Recorder 1 may have internal storage inside the server while Recorder 2 requires an external direct attached storage (DAS) device. Every server is a physical machine requiring memory, CPU, power, cooling and rack space in the data center. This could be a pretty costly proposition. Some questions may arise, such as: What would happen if Recorder 1 would crash or fail? What would happen to all the video data that has been stored on it? What would happen if both recorders would run out of storage space and more cameras would be needed?

A system that is improperly designed could have significant effect on the overall surveillance environment. Frame loss, difficulties to scale, and slow performance are just a few effects of an improperly designed and misconfigured video surveillance system.

Part of Hitachi Visualization solutions, Hitachi Visualization Platform (HVP) offers a choice of video management platforms. The HVP video management platforms take a shot at solving all the problems of having a proper video surveillance infrastructure. These converged solution appliance solutions comprise storage and compute hardware that run digital video management system software. Taking scalability, high availability, redundancy, data center space and cost into account, HVP video management platforms leverage Hitachi Data Systems hardware with VMware hypervisor to build an amazing environment that could easily sustain small to large-scale video surveillance deployments.

HVP video management platforms are designed around multiple video management systems and support all the major VMS companies in today's market. This article focuses on the complexities of video management system environments and how the HVP video management platforms solve these challenges.

HVP Video Management Platform Offerings

Hitachi currently offers three models of the HVP video management platform: HVP 150 video management platform, HVP 500 video management platform and HVP 1000 video management platform. All three systems use Hitachi Compute Rack 220H hosts and Hitachi enterprise storage systems such as Hitachi Unified Storage (HUS) and Hitachi Virtual Storage Platform (VSP), a choice of Brocade or Cisco MDS switch fabrics for backend storage networks, and VMware for managing the overall environment. Below are the specifications currently for every model we support.

HVP 150 Video Management Platform

- Supports 150 2.0 megapixel cameras at 30 FPS, H.264 compression.**, ***
- 2x physical hosts (2x Intel Xeon E5-2620 6-core 2.0GHz with hyperthreading) per host.*
- Either (4x 1Gb NICs and 2x 10Gb) or (8x 1Gb NICs) per host.
- 32GB of DDR3 RAM per host.
- 2x 300GB 10K SAS drives.
- Brocade 320 or Cisco MDS 9148 switch fabric.
- HUS 130, supporting a minimum of 30TB to a maximum of around 800GB of storage.
- 4x virtual machines (VM) (VMware vCenter Server Foundation, Management Server, 2x Archivers).
- Upgradable to 300 cameras with additional host and storage.

HVP 500 Video Management Platform

- Supports 500 2.0 megapixel cameras at 30 FPS, H.264 compression.**, ***
- 2x physical hosts (2x Xeon E5-2625 8-core 2.6GHz with hyperthreading) per host.*
- Either (4x 1Gb NICs and 2x 10Gb) or (8x 1Gb NICs) per host.
- 64GB of DDR3 RAM per host.
- 2x 300GB 10K SAS drives.
- Brocade 320 or Cisco MDS 9148 switch fabric.
- HUS 150 supporting minimum of 30TB to a maximum of around 2.4PB of storage.

- 7x VMs (vCenter Server Foundation, Management Server, 5x Archivers).
- Upgradable to 1000 cameras with additional host and storage

HVP 1000 Video Management Platform

- Supports 1000 2.0 megapixel cameras at 30 FPS, H.264 compression.**, ***
- 3x physical hosts (2x Xeon E5-2690 8-core 2.9GHz with hyperthreading) per host.*
- Either (4x 1Gb NICs and 2x 10Gb) or (8x 1Gb NICs) per host.
- 128GB of DDR3 RAM per host.
- 2x 300GB 10K SAS drives.
- Brocade 6505 or Cisco MDS 9148 switch fabric.
- VSP G600 supporting a minimum of 30TB or a maximum of around 20PB of storage.
- Upgradable with add-on hosts and storage.
- 13x VMs (vCenter Server Standard, Management Server, 10x Archivers, 1 Microsoft® SQL Server®).
- Upgradable to nearly 10,000 cameras with additional hosts and storage.

* Additional servers may be used for LPR cameras, video analytics, access control and any other VMs pertaining to video infrastructure including Hitachi Visualization Suite.

** Higher resolution cameras are supported with lower camera quantities per each solution.

*** Lower resolution cameras are supported with less physical host resources, allowing more cameras to be added to physical host.

Video Surveillance History

Not too long ago, casino control rooms were filled with VCRs utilizing VHS tapes to record video. Miles upon miles of coax cable would be laid for the camera connections. There were hundreds of Cathode Ray Tube CRT monitors inside operations centers connected to every single feed. Most monitors would support one feed from the camera while some would support multiple windows on the screen. Features such as analytics, user profiles and multiple frame rates did not exist. Video traffic would never interfere with IT network traffic. Audiovisual (AV) engineers and technicians who had no knowledge about IP networks, servers or storage systems would run video infrastructure.

As the world pushed towards the digital age, so did the video surveillance market. A few years later, cameras were no longer recording to VHS tapes, and engineers had to completely forget what VHS tapes were. Our cameras required IP addresses, and words like “managed switches,” “firewalls,” and “routing” were slowly introduced into the new world of video surveillance. IT traffic and video traffic no longer had to be physically separated and they started to share the physical infrastructure and bandwidth. Acronyms and words such as “QoS,” “VLANs,” and “multicasting” resonated when building out proper video environments. Servers with DAS, NAS and SAN storage attachments started replacing VCRs while digital monitors replaced CRT monitors. Network and systems engineers now maintain video surveillance environments and electricians maintain the physical camera installations and repair.

Figure 1 perfectly illustrates the next phase of video surveillance architecture. As cameras became digital, new software companies emerged that provide a way to manage these cameras. Today, Milestone, OnSSI, Genetec, Verint and NICE Systems are some of the key players in the VMS market. These VMS companies mostly have the same master server and archiver architecture. All of these video management systems run on Microsoft Windows® operating system. Users have client software installed on their local machines to allow them to connect into the server environments and control their surveillance system.

Figure 2 illustrates Genetec's Security Center 5.1 client. The user is able to see groups of cameras on the left and drag them into viewing panes on the right. The large blue circle on the bottom-right allows the user to have pan/tilt/zoom capabilities as long as the cameras support these functions.

Figure 2. Genetec Security Center Client

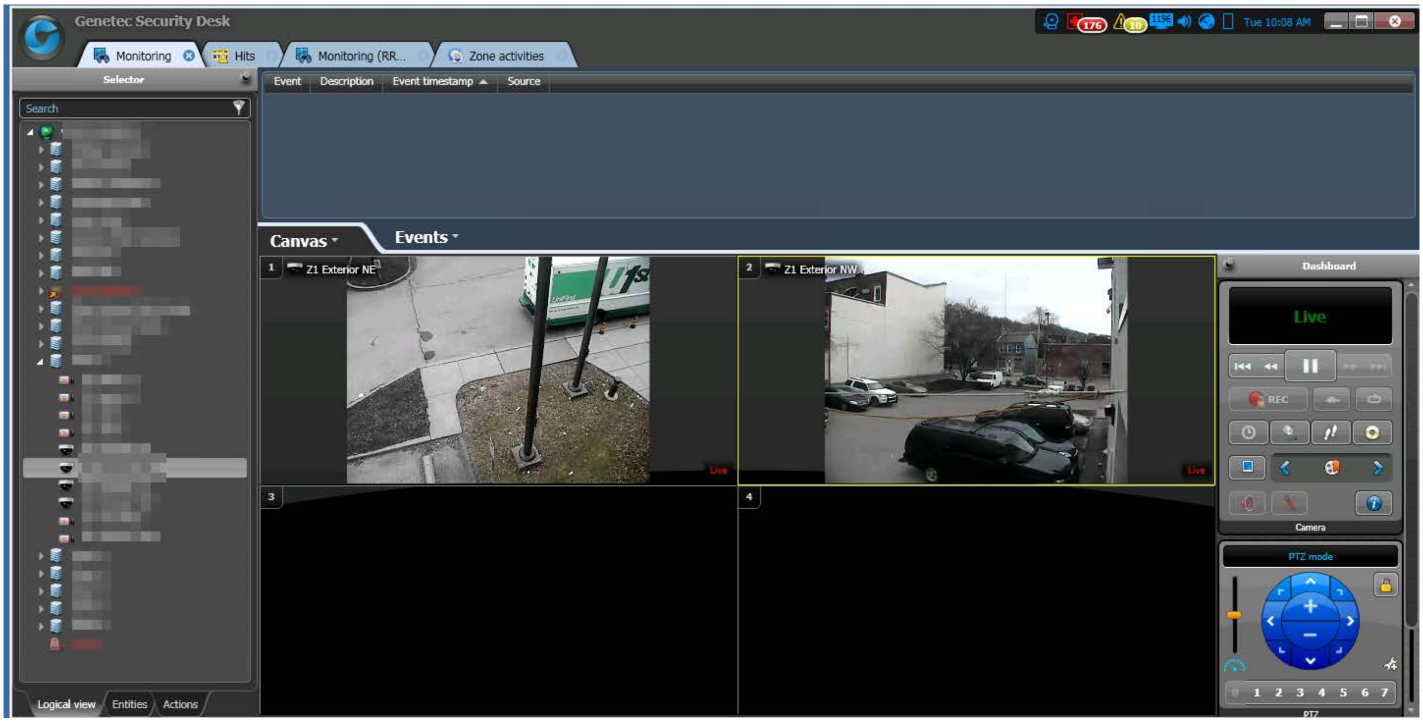

The new model for designing video surveillance environments involves virtualization and cloud computing. Virtualization adds transparency and a new hypervisor (virtualization) layer to the overall environment. With introduction of virtualization, the new environment looks very similar to Figure 3. From the looks of it, things may seem a bit confusing, but we will tackle VMware's role in the video system architecture later in this white paper. No longer are the archivers or the master servers bound by physical locations or CPU and memory constraints. The VMS is installed onto virtual machines that are easier to maintain, support and scale. Video can now be recorded on-premises or somewhere in the cloud where it can be retrieved anywhere in the world. This is the new direction for not only IT but also video surveillance architectures. By utilizing access memory and CPU from physical server hosts, we can now save on power and cooling costs, as well as physical space in the data center. The time of deployment of new infrastructure shrinks from weeks to minutes, which reduces overall business costs associated with procurement and installation of physical hardware. This is the new age of computing, yet some might see this as stepping back into the days of mainframes.

Figure 3. Hybrid Cloud Architecture for Hitachi Visualization

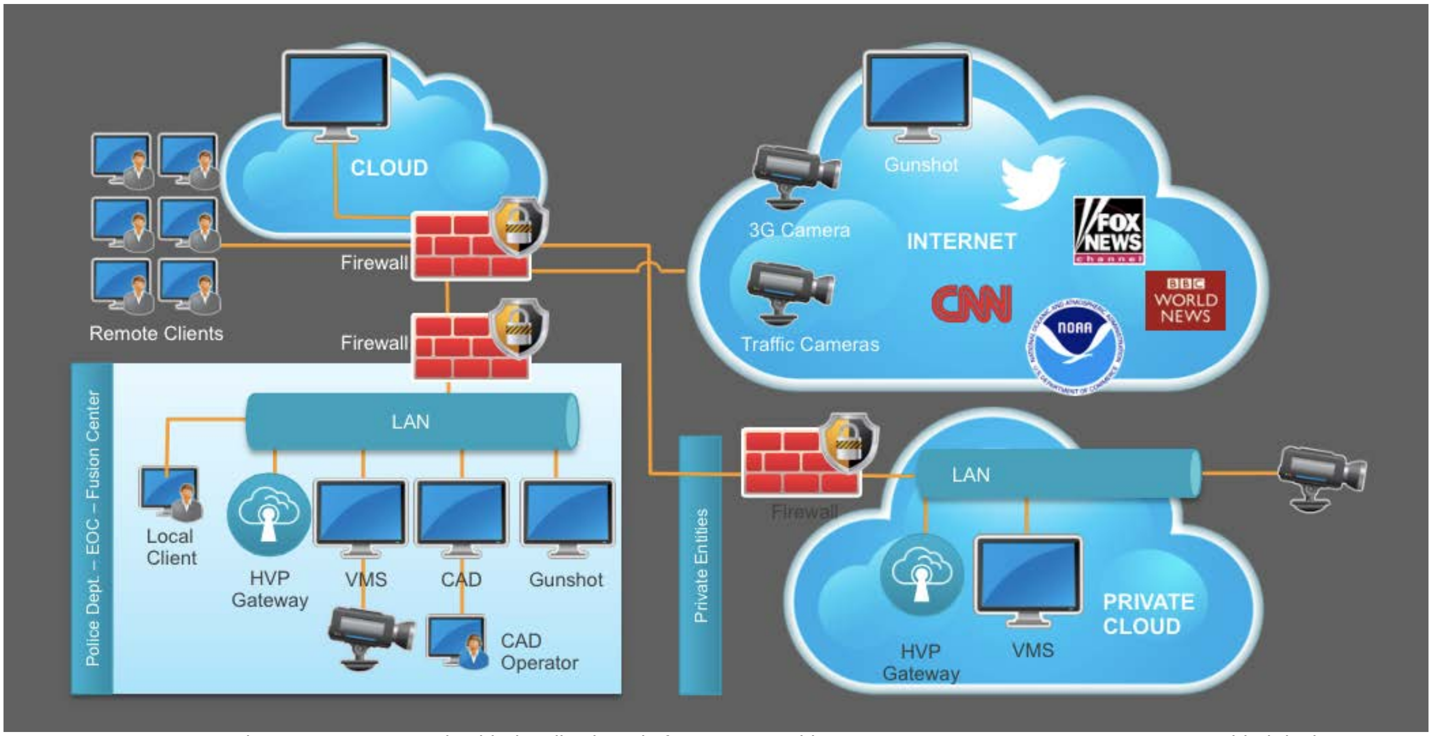

EOC = emergency operations center, HVP = Hitachi Visualization Platform, VMS = video management system, CAD = computer aided device

Hitachi Visualization Platform: Video Management Platform Architecture

If video surveillance environments were designed and treated the same way as an enterprise IT solution for databases, file storage, performance degradation and frame loss would become noticeable over time. Even though most IT infrastructure principals apply, video surveillance systems have specific caveats that must not be overlooked. An improper video surveillance infrastructure could be costly in the long run since installing additional hardware would be the cheapest way to compensate for basic requirements. A great example of this would be adding storage because the slower hard drives cannot handle the IO throughput. The system would also be too tough to scale in a cost-effective manner. It is important to focus on numerous aspects of a video surveillance environment to engineer a proper solution.

The HVP video management platform design focuses on numerous aspects of the video surveillance architecture. Five main directives (pillars) drive the overall design of HVP video management platform: **storage, compute, throughput, scale** and **cost** with VMware and redundancy injected into every pillar. This paper considers the complexities of each of the five topics and how HVP video management platform addresses each of them. It is important to know that these complexities correlate with one another. For instance, not having a scalable system will drive the cost higher when trying to expand the infrastructure; or improper connections to storage will drive CPU utilization up. This paper explains how IT best practices and proper design drives the approach for HVP video management platform it is considered the best video management system appliance on the market today.

Figure 4. VMP Design for Genetec and Verint

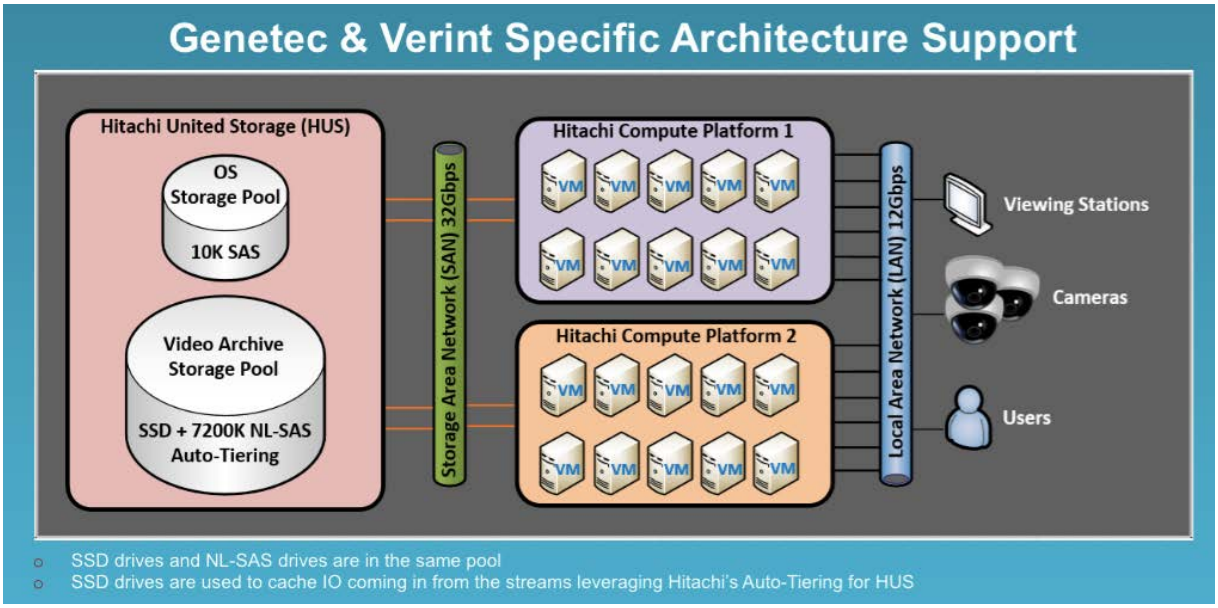

VM = virtual machine, SSD = solid sate disk, NL-SAS = nearline serial attached SCSI

Figures 4 and 5 illustrate a high-level overview of the HVP video management platform design. It is important to present the high overview before getting into the nitty-gritty of every aspect of this platform. Ahead, this paper considers every component of the designs shown in Figures 4 and 5.

Figure 5. VMP Design for OnSSI and Milestone

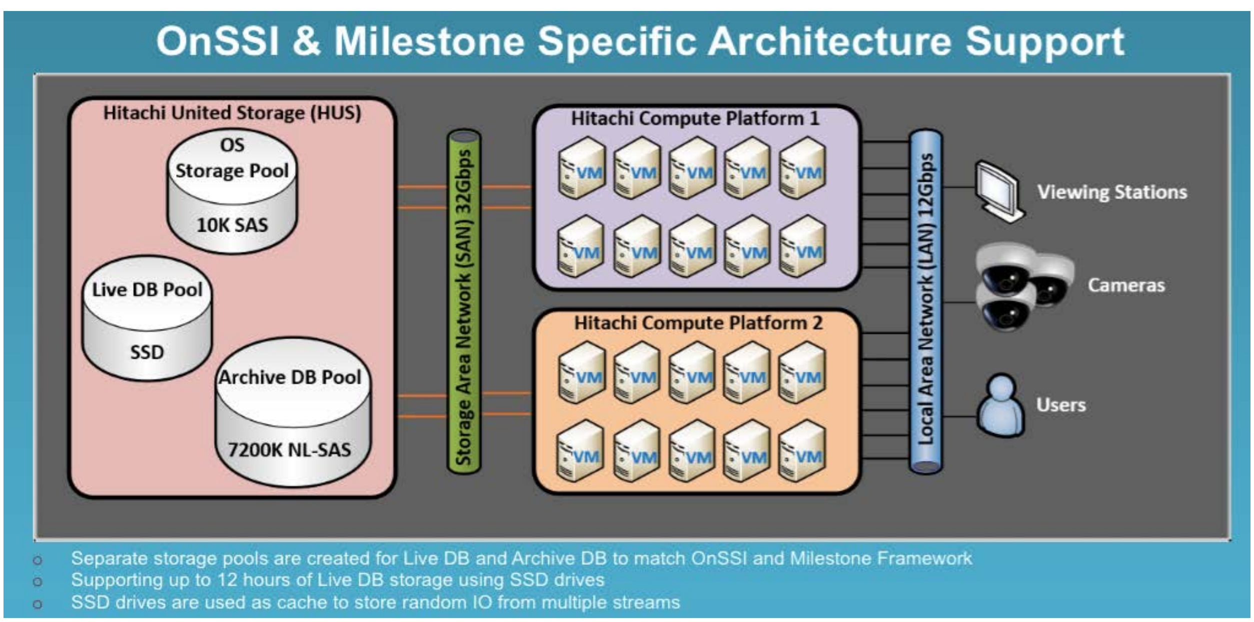

OS = operating system, DB = database, NL-SAS = nearline serial-attached SCSI, VM = virtual machine, IO = I/O

VMware

VMware plays a central role in HVP video management platform since it's the driving force of solving high-availability, redundancy, scalability and cost issues of the video surveillance architecture. Once again, it is injected in every level of the five pillars and it is impossible to talk about them without mentioning or discussing virtualization in every one of the topics below. Without VMware, every recorder would have to be a physical server that would take up extra space in the data center, requiring extra power and cooling solutions. These servers would be hard to scale. Redundancy would be relying on VMS failover recorders or failover management servers also requiring physical hardware. VMware controls every aspect of HVP video management platform, making it one of the most important components in the overall solution.

Compute

Gordon Moore, one of the founders of Intel, predicted that the number of transistors on a CPU would double every year and that this trend would continue to this day. Moore's law has proven to be pretty accurate, since CPU power nearly doubles annually as he predicted. Every year, our home computers, laptops, servers and even cell phones become more powerful. In the video surveillance world, this is relevant. Every year, the market produces more powerful cameras with better resolution, which requires more CPU power on the recorders to handle the higher resolution. Most environments typically do not support more than two megapixel cameras. This reality creates an interesting trend. As CPU power increases on an annual bases the CPU requirements for video surveillance architecture doesn't change much necessarily. A 2 megapixel camera is relative to 1600 x 1200 pixels, which is significant video quality. Because most video management systems don't recommend more than 150 camera connections per recorder, we can now bring virtualization into play and virtualize our recorders to utilize CPU and memory of our physical hosts more effectively. Referencing Figures 4 and 5, we see that multiple virtual recorders are now part of only two physical hosts with enough power to run them. Recording video does not really require much memory or CPU utilization. Playing video clips, encoding or decoding video data does require higher CPU utilization. Unless we need to decode video surveillance footage for one reason or another, we don't need to have multiple physical hosts in the architecture.

HVP video management platform architecture supports two physical hosts out of the box for HVP 150 video management platform and HVP 500 video management platform and three physical hosts for HVP 1000 video management platform. The system can tolerate an entire host failure and would still be able to function. HVP video management platform has full redundancy on the host side, including multiple network interface cards (NICs), Fibre Channel host bus adapters (HBAs), connections to both storage and LAN networks, power, host drives running ESXi, and hosts themselves. The system may drop an entire host and still be able to recover.

Network Throughput

It is incredibly important to determine how much throughput a system can handle. We base this on ingress and egress traffic. Ingress traffic is mainly from all the camera traffic coming into the HVP video management platform while egress traffic are the users pulling video data from the system. Not only does the HVP video management platform provide enough throughput for video traffic, but it supports out-of-band management traffic as well. Each hosts has either (6x 1Gb NICs and 2x 10Gb) or (8x 1Gb NICs). With this many NICs, we encourage organizations to double up on their LAN side switching capabilities. We always dedicate one physical NIC for management and one for VMotion traffic (VMs moving from one physical host to another) for VM migration during maintenance periods. The rest of the adapters are used for video traffic.

Two or three megapixel cameras may sometimes produce 1Mb/sec to 10Mb/sec per camera stream. It's important to make sure that the LAN side of HVP video management platform can handle that throughput. It's also important to note that some environments strictly do not need and cannot handle megapixel cameras. Their network might not be able to support or sustain the throughput requirements. No matter what the camera bandwidth requirements are, HVP video management platform is designed to handle large amounts of throughput.

Storage

Storage is a very important topic in the video surveillance architecture and one that is very easily overlooked. In order to discuss storage, we must first understand the complexities of video I/O and how we must properly record it to disk. Table 1 sums up the complexities of video I/O and we examine each item, one by one.

As we discuss these storage complexities, we must also explain how HVP video management platform solves and embraces the complexities of storage. HVP video management platform architecture is built around understanding the video storage environment.

Table 1. Video I/O Explained

Large, Random I/O	▪ Multiple cameras attach to an archiver, each with its own resolution, stream and data. ▪ Cameras pointed at a door will need less bandwidth than a camera looking at a busy intersection.
95% Write, 5% Reads: Reads Are Spiky	▪ Video recording is constant. Thus files are writing to storage all the time. ▪ Video is viewed mainly after an event occurs, and multiple users try to pull the same streams.
Video Traffic Is Fragmentation Prone	▪ As video is groomed off after a certain time period, fragmentation occurs in the drive, which may decrease performance.
Different Digital Video Management System (DVMS) Vendors Use Different Methods for Archiving	▪ Milestone and OnSSI support storage tiering. This helps improve random write I/O when it's moved from live to archived database (requires as separate storage pool). ▪ Genetic and Verint need caching. This allows higher-end drives to distribute I/O to the slower drives using auto-tiering.
Most DVMS Vendors Recommend Using RAID-6 for Redundancy	▪ N+2 redundancy for video storage. ▪ Performance decreases as redundancy increases.
Disk Throughput Is Just As Important As Capacity	▪ Using 256KB block size instead of 64KB increased throughput.
Fibre Channel Host Bus Adapter (HBA) Solution Reduces CPU Utilization on the Host	▪ Offloading storage I/O processing to HBAs reduces CPU utilization by 10%.

LARGE, RANDOM I/O

Video I/O is large and random in nature.

You would think that video traffic is large, sequential; however, this is not true. A recorder with over 100 camera connections records data from every single one of these data streams. Each camera may be configured with different frame rate and encoding standard, and may literally be pointed at different environments.

A camera pointed at a door with not much in the area of pixel changes will produce more traffic than a camera looking at a very busy intersection in the street. Although, this white paper does delve deeply into compression mechanisms, we do want to note that these mechanisms may help drive these differences.

95% WRITE, 5% READ: READS ARE SPIKY

Mostly, video is continuously being recorded to video management system or VMS software. Video is constantly being recorded to the drives, which is why we write data 95% of the time. What most administrators don't know is that when a user tries to pull a live video image, they seldom ever pull it from a camera. They always pull it from storage, and it may milliseconds off real time. Imagine an event that causes 20 people to start looking at live video simultaneously. Thus, 5% of the time we will produce reads. Reads mainly occur during an event occurrence, which causes everyone to pull up live video at the same time, which then causes the reads to be spiky.

VIDEO TRAFFIC IS FRAGMENTATION-PRONE

Because video I/O is random, it will be scattered around the disk platters as disks keep spinning to record data. When the older video is groomed off (erased) to make room for newer video files, the data will be removed from different parts of the disk, leaving chunks of holes inside it. Disks don’t like scattered data. It’s much easier to seek data located in the same memory space of a disk than in different locations. It’s important that I/O written to slower disks is contiguous and same data is easier to find because it will reduce frame loss over time.

DIFFERENT DVMS VENDORS USE DIFERENT METHODS FOR ARCHIVING

Some VMS vendors understand that to write data contiguously to disk as discussed previously, data must first be written to faster disk that can handle the I/O. Then, that data can be offloaded to a different slower disk volume. OnSSI and Milestone both have special “Media Database” or “Live Database” volumes that are composed of faster disks. These volumes are considered first tiers of storage from which video gets viewed or tiered to slower disk volumes over specific periods of time.

Unlike OnSSI and Milestone, other VMS vendors such as Verint and Genetec do not have that capability. They rely on proper engineering of the video architecture environment to either build enough of a caching mechanism to sustain the I/O or build auto-tiering into their solution.

MOST DVMS VENDORS RECOMMEND USING RAID-6 FOR REDUNDANCY

Many video appliances are installed and configured by systems integrators instead of internal IT. The mindset that video surveillance infrastructure is not an IT asset is still the approach as this document is being written. Because of this, most VMS companies expect that even though RAID-6 is very write-penalty heavy, it should still be used as the main striping mechanism for video management systems. RAID-5 does not provide enough redundancy and RAID-1+0 could be very expensive when discussing 200TB or 2PT of required storage.

DISK THROUGHPUT IS JUST AS IMPORTANT AS CAPACITY

Figure 6 is an actual screenshot to calculate disk requirements for 560TB of usable storage capacity. The diagram shows the differences between using 64K block size versus 256K block size. With 64K block size you would need 650 disks to achieve 750MB of throughput, where with 256K block size you would need 200 drives to achieve the same capacity. This clearly demonstrates that throughput planning is just as important as capacity planning.

Figure 6. Disk Requirements Calculator

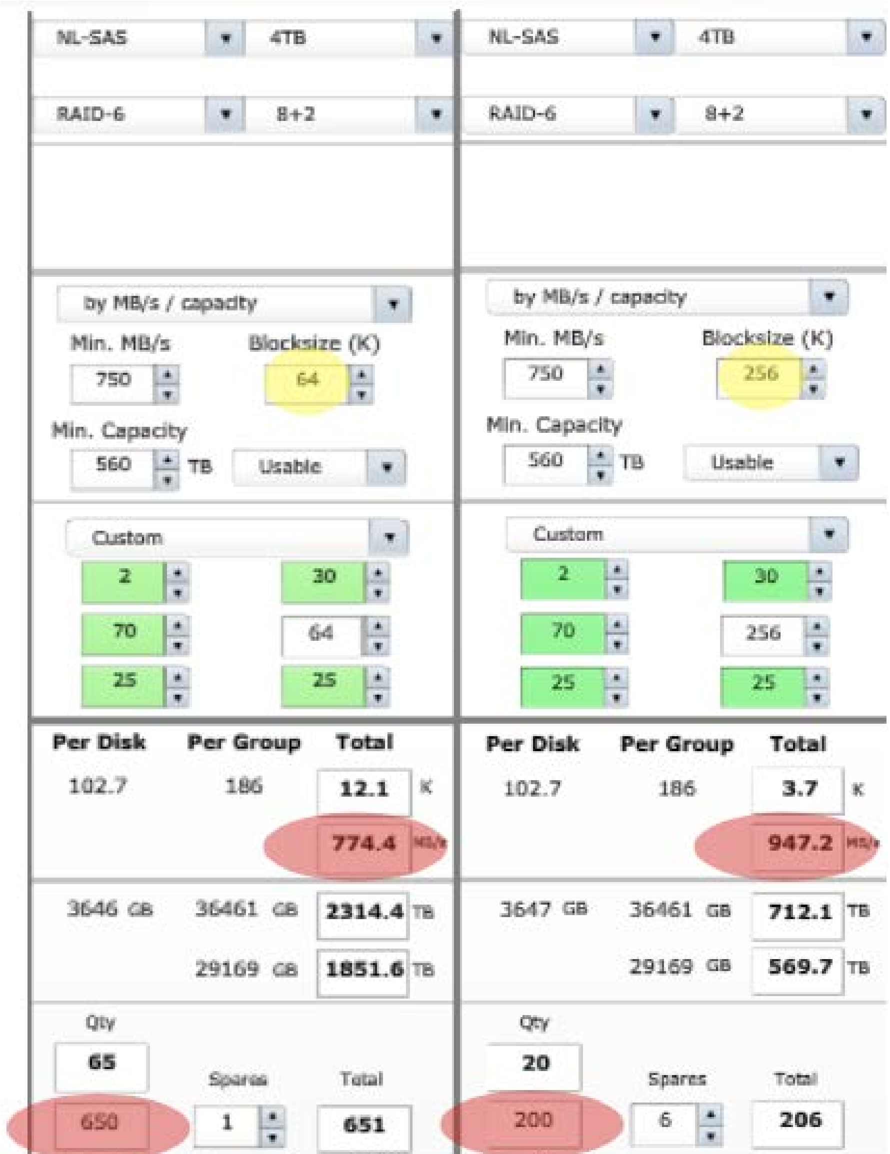

Based on HiPace Calculator

FIBRE CHANNEL HBA SOLUTION REDUCES CPU UTILIZATION ON THE HOSTS

It is becoming common knowledge that iSCSI has become a much utilized storage protocol in the industry. There are definite benefits to use one instead of the other. iSCSI runs over existing LAN infrastructure on top of TCP/IP. Fibre Channel is specific to storage and requires its own specialized SAN fabric switches. Two very important concepts are important to grasp: First, Fibre Channel is specifically designed for storage networks and has a much better protocol since it makes sure that the host is ready to receive I/O traffic before it is sent to the host. Second, using HBAs reduces CPU load on the physical hosts by 10%.

HOW MUCH STORAGE IS ENOUGH?

Storage capacity is dependent on a number of factors, and determining the exact amount needed may prove to be difficult at times. Camera resolution, frame rate, compression mechanism, motion in the camera image and even the camera, itself, play a part in storage requirements. If an environment is only recording on motion and 90% of the cameras are looking at doorways with still images, then storage requirements could be greatly reduce. There are online calculators that help determine storage capacity; however, because of the complexities in storage environments, they are not always accurate. Some calculators may present different storage requirements based on the same exact data inputs. Just to understand approximately how much storage is required we look to Figure 7. The chart shows 1x 2 megapixel camera recording for just a single day at 30 FPS may need around 37GB for storage, while 100 cameras recording for 30 days with the same resolution and frame rate may need close to 111TB.

Figure 7. Storage Calculator

Number of channels	Type	Manufacturer	Camera	Activity	% Recording	Days	FPS	Bitrate (Kbps)	Storage (GB)
1	IP	Axis	Q6045 Resolution: 1080p Compression: H.264	high	30 Audio	1	30	11500 Custom	37.26
1	IP	Axis	Q6045 Resolution: 1080p Compression: H.264	high	30 Audio	30	30	11500 Custom	1,117.80
100	IP	Axis	Q6045 Resolution: 1080p Compression: H.264	high	30 Audio	1	30	11500 Custom	3,726.00
100	IP	Axis	Q6045 Resolution: 1080p Compression: H.264	high	30 Audio	30	30	11500 Custom	111,780.00

Hitachi takes into account all the storage complexities, and builds its HVP video management platform appliances around these factors. This is what really separates this appliance from regular IT infrastructure solutions or competitors in the industry.

HVP video management platform records data to two different tiers of storage. The first tier is composed of faster disks that are capable of handling large amounts of write and read I/O. These disks support live viewing capabilities without disrupting constant video writes to the disk. These disks are expensive: Thus, we tier the video to the second tier of cheaper and larger NL-SAS disks. When we move data from faster tier drives to NL-SAS the data is moved to the slower drives contiguously as it is also groomed off the faster drives contiguously. Once the data is erased from the drive, the pool reclaims the empty space. This solves fragmentation issues on the disks.

For Genetec and Milestone we enable the auto-tiering feature and create one pool volume comprising faster and slower disk. For OnSSI and Milestone, for instance, we create three pools of storage and have the HVP video management platform tier the video itself. Figures 4 and 5 show both options.

Both live database and archive volumes use RAID-6 since using faster disks in the first tier of storage allows HVP video management platform to take the RAID hit without noticeable performance decreases. That said, both tiers of storage support 256K block size to account for the I/O throughput requirements.

HVP video management platform uses the Fibre Channel SAN connectivity, since it reduces the CPU load of the host and is considered to be a better storage protocol overall.

Scale

HVP video management platform is designed to be a scalable solution. HVP 150 video management platform supports 150 cameras out of the box with a minimum of 30TB of storage. It can scale to 300 cameras by addition a third host and can scale up to 2PB of storage as needed.

HVP 500 video management platform can scale to 1000 cameras by adding an additional host. And, it can scale from 30TB of storage to 2.5PB of storage as needed.

HVP 1000 video management platform can scale up to 10,000 cameras and up to 2.5PB of storage.

Cost

HVP video management platform is designed to be a cost-effective solution without sacrificing most important aspects of the appliance, such as performance and redundancy. By separating storage from compute we allow organizations to buy either storage or compute an as-needed basis. They can avoid stacking servers for additional storage without a current need for compute and memory power.

Building a second tier for larger, cheaper disks guarantees that an organization will not need 15K SAS or SSD disks throughout its entire environment.

Using only two or three hosts in the environment reduces the costs of power and cooling in the data center as well as rack space that is needed strictly for storage.

Being able to deploy additional recorders without adding physical installations of new hardware saves deployment time. It eliminates hardware identification, procurement, racking, stacking and most configurations, reducing overall operational costs.

Conclusion

The Hitachi Visualization Platform offers video management platform architecture that is designed by engineers with real-world experience in systems integrations and working with organizations in the video surveillance industry. It is a product that came out of real-world necessity. It solves real-world challenges of administering ever-increasing amounts of video in storage environments. Focusing on the five main pillars (storage, compute, network throughput, scale and cost), it is considered a one-of-a-kind solution in the market today.

The HVP video management platform is the first true highly scalable and available, fault-tolerant, cost-optimized solution for video processing and management. It is currently going through certifications with the major VMS vendors to ensure overall performance and to validate each VMS environment through proper testing in a lab environment.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 95050-2639 USA
www.HDS.com **community.HDS.com**

Regional Contact Information
Americas: 1 866 374 5822 or mailto:info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

WP-536-A DG July 2015