S O L U T I O N   P R O F I L E

Five Criteria for Video Surveillance Storage

June 2007

THE NET NET

Evolving far beyond its security guard monitoring roots, video data has become a key strategic asset to many enterprises, and is poised to change business practices for many others. Much of this change has been brought about by an evolution of technology that allows video to be delivered and managed as a digital asset rather than purely a real-time observation tool. Those changes have included several innovations:

- Encoding technologies that digitize and highly compress video, making it possible to transmit video in new ways and store video on different types of media.
- Prolific TCP/IP transmission mediums that have allowed new encoding technologies to use cheap, easily installed, and robust data connectivity technology for cameras.
- Low cost, high capacity, high performance storage that allows compressed video to be stored and repeatedly accessed.

Nonetheless, integrating all of the elements of a digital video solution is complex and carries significant consequences. One of the most significant elements of a video solution is the storage system. In fact, over 40% of a digital video solution is spent on storage. Insufficient storage can easily render a sophisticated video system useless. Further exacerbating the problem, traditional storage systems that typically service files and databases are by no means easy to configure with the performance, scalability, and availability necessary to support video.

In this solution profile, we first review how digital video is driving the need for more capable storage, the different types of storage architectures available today, and the five key requirements to consider when evaluating a storage solution for video storage. Lastly, we spotlight Pivot3 as a vendor focused on delivering a high performance, cost-effective, easy-to-use storage system for the demands of digital video and how their solution compares with the core video storage requirements that we set forth.

Video and Evolving Technology

Digitization has changed the value of video. In the past, enterprises deployed analog video technologies through closed circuit television (CCTV), where video channels were streamed to a central multiplexer over coaxial cable run within a building. The multiplexer interleaved multiple channels of video onto a tape, shifting between each channel at different time intervals. As a result, CCTV compromised the rate of capture in order to support more cameras while using less tape. In the end, video was not highly useable because tapes were awkward to access, and the granularity of capture was insufficient to see needed detail.

Today, video images are digitized and stored as data files with all of the benefits that digital systems offer over their analog counterparts. Specifically, digital video can now be easily searched, quickly accessed, computer-analyzed for embedded content and then distributed for wide-spread use. These attributes of digital video are driving the replacement of traditional CCTV systems with digital video systems.

The first generation of digital video began with isolated Digital Video Recorders (DVRs) that captured and digitized traditional analog video. Today, this has evolved into Network Video Recorder (NVR) systems with distributed in-camera processors that create digitized video before transmitting it over data networks for storage.

The Challenge of Video Storage

Concomitant with the deployment of digital video systems, the amount of generated digital video data has grown exponentially as more cameras with higher resolution and faster frame rates are deployed and video data is kept longer to meet heightened security risks.

The end result is a massive increase in the data storage requirements for digital video. Higher capacity points, greater availability demands, centralized management and the need to distribute data over an IP network are all leading end users to rethink their video storage infrastructure.

Traditionally, video managers have had to choose among three different storage approaches to meet this data onslaught. However, each of these traditional approaches suffers its own drawbacks. Let's look at each of these architectures and its challenges:

Direct Attached Storage (DAS)

DAS relies on a storage controller within a server or specialized cabinet that is attached to a server. The storage controller virtualizes a group of drives together into a RAID set (see sidebar on RAID), and presents the disks as a single storage volume. DAS is an economical and simple solution for low capacity, low performance environments. However, to add large amounts of capacity, multiple DAS servers must be deployed, each with its own management and storage provisioning overhead. Therefore, with increased capacity needs, the overall management of a DAS environment becomes cumbersome.

Furthermore, DAS servers quickly become performance bottlenecks, and even with redundant hardware can become a single point of failure from routine configuration changes, security issues, or operating system patching. The single point of failure of DAS also becomes cumbersome during capacity additions or routine operational changes that result in system downtime and can render the entire video system temporarily inoperable.

Given the 24x7 nature of video capture and surveillance, downtime whether planned or unplanned can result in the unacceptable loss of critical video data. Moreover, in regulated environments (e.g. casinos) or high security risk areas (e.g. nuclear power plants), the loss of video data can result in significant monetary losses and safety and compliance concerns.

Network Attached Storage (NAS)
NAS systems present storage over a traditional IP network using standard file access protocols like NFS or CIFS. Typically, NAS systems are dedicated, turn-key, highly-available, and hardened appliances. However, because these systems are appliances, they can only be scaled to preset performance and capacity levels without requiring a forklift upgrade to a more expensive, higher performance system. These scalability limitations force video managers to predict their long-term requirements and buy capacity and performance for the future up front, usually expending many more capital dollars than justified by initial demand. If a sizing error is made, or a system grows beyond initial expectations, costly total replacement and complex data migration may be required.

Fibre Channel Storage Area Network (FC SAN)
FC SANs attach standalone storage arrays to one or more servers over a specialized fiber network. A FC SAN provides a proven mechanism to scale capacity and deliver high I/O performance, but these networks add significantly more complexity and cost than NAS or DAS to purchase and operate. Further, FC SANs require highly specialized skill sets that are rarely found outside of enterprise datacenter environments. Given that video surveillance and capture infrastructure is increasingly becoming IP-based, introducing another network topology runs counter to the rest of the video surveillance and capture infrastructure (e.g.

What is RAID?

RAID stands for Redundant Array of Independent Disks. While RAID technologies have a number of applications, the most common feature in storage systems is protecting stored data and/or increasing performance. RAID techniques are applied to an individual storage volume that is composed of multiple physical spinning disks. Several variations of RAID can protect the volume in different ways. These include:

RAID 1: Creates duplicate copies of the entire data set on another similar disk. This delivers high levels of protection, but doubles the amount of storage required to store the data set.

RAID 5: Creates an extra parity bit on a single disk calculated from all of the other bits distributed across other disks in a storage volume. In a collection of disks, the capacity of one can effectively be used for parity bits. A failed disk can be replaced and the data re-built from the remaining data and parity bits. This requires only one more disk than raw data, thereby optimizing capacity, but writing a parity impacts performance.

RAID 6: Creates double the parity of RAID 5, allowing two disks to fail, but uses two more disks than raw data alone, while requiring more calculations and disk overhead.

The occurrence of disk failures is routine in active storage systems. Without RAID, data loss will occur, and is more likely to occur in larger storage systems that use more disks.

IP cameras, IP access control and intrusion detection devices and next-generation bio-credentialing systems) needed

A New Storage Choice for Video – IP-based Clustered Storage

Until recently, video surveillance systems had no other choices beyond these recognized storage approaches. Today, IP-based clustered storage presents a cost-effective alternative that meets the rigorous demands of video storage. Clustered storage can aggregate multiple distinct, affordable storage nodes into a single, easy to manage, high performance storage cluster that is accessible over Ethernet network using the iSCSI protocol. Clustered storage can be cost competitive with DAS while overcoming the performance and scalability limitations of all of the other architectures. Altogether, clustered storage is well suited for the high performance, high capacity scale of video storage environments.

Criteria for Evaluating Video Storage

Next, we posit five criteria that end users should consider when selecting between DAS, NAS, FC SAN, and IP-based clustered storage options for their video surveillance deployments.

Criteria #1: Networked Storage

Storage networked over Ethernet or Fibre Channel provides several important capabilities. First, networked storage increases system versatility by allowing interconnection between multiple systems in many locations. This allows data to be stored from widely distributed cameras, and controlled access by workstations, management applications, or backup systems for managing, sharing, and protecting video data. Second, just as with traditional NAS systems, networked storage systems provide simplified administration with hardware purpose built for availability and performance.

Criteria #2: Easy to Use, Easy to Manage

Ease of use in a storage system is often overlooked by business managers responsible for video. But without an easy to use interface, managing storage often requires specialized and high-cost skill sets. Common video storage systems may seem simple, but prove highly complex to administer when a disk failure occurs or storage needs to be expanded. A storage system should provide easy, non-disruptive administration for all tasks.

Criteria #3: Fault Tolerance

Availability and reliability are indicators of how resilient a storage system is to the loss of function due to a storage component failure or to the loss of data due to a corruption. Most storage administrators will encounter one or both of these failures during the life of a complex storage system. Redundancy is crucial to mitigate failures. This includes redundant components that take over operations during a failure and sets of multiple disks that are configured to protect data.

Common video storage systems today utilize redundant components, but may be configured behind a server that becomes a single point of failure. What is worse is the real possibility of a server failure due to mis-configuration or routine patch and maintenance activities. In addition, large standalone storage arrays may be subject to outages during software upgrades and patches. Ideally, a storage system should be able to survive multiple failures at the controller and disk levels and continue to process data without interruption.

Criteria #4: Expandable Capacity
Expandability within a single system is crucial for video storage. With the increasing value we see in captured video today, we believe that most organizations will retain and reuse more video in the future. Expanding storage to keep up with the growing demand to keep more video data online longer can be a costly and time consuming proposition, requiring professional services and disrupting video capture. Traditional storage systems are limited in expansion headroom, and require significant, up-front investments to buy extra capacity and performance in order to meet uncertain future needs. An ideal storage system should scale gradually and non-disruptively without the need to invest in expensive additional controllers or unneeded capacity, allowing an enterprise to pay for disk and performance as required.

Criteria #5: Scalable Performance
Scalable performance is the most complex aspect of video storage, and is one of its most crucial requirements. As video usage expands, I/O performance demands often exceed the capabilities of the underlying storage subsystem. Yet expanding storage goes far beyond simply adding additional drives and capacity. In many storage systems, expanding performance requires a complete replacement of key components or even the entire storage system.

Video requires a cutting edge storage design that allows administrators to efficiently add processing capabilities, additional disk, ports, and throughput to a storage system, while preserving performance. Traditional storage systems do not scale well in performance, and incur significant downtime when hardware upgrades for scalability are undertaken. Initially, storage systems are purchased with high performance controllers, but those same controllers become bottlenecks as more disks and clients are added. Many systems can expand cache to burst at high speed for a short period of time, but this is insufficient for video and the demands for long sustained performance.

Focus on Pivot3

Pivot3 has developed an innovative IP-based clustered storage architecture ideally suited to the rigorous storage requirements of video. Pivot3 delivers optimal price-performance, flexible scalability, and fault tolerance in an easy to use and deploy system. Pivot3's clustered storage allows the performance and storage capacity of the system to be scaled independently to meet the needs of a particular application, while maintaining one of the most cost-effective price points on the market.

The heart of Pivot3's storage is a clustered architecture that utilizes industry standard hardware running a proprietary Pivot3 RAIGE (RAID Across Independent Gigabit

Ethernet) operating system. RAIGE is a simplified, optimized, embedded Linux kernel that turns a server into a highly available, purpose-built storage appliance.

The RAIGE OS delivers a wealth of availability features on top of industry standard server hardware. RAIGE protects each drive in a 12-drive appliance with tunable volume-level parity or mirroring choices, effectively protecting against any single drive loss. Furthermore, multiple storage servers (called Databanks in Pivot3 terminology) can be added together into a storage cluster that additionally protects data across the Databanks. Databanks can be added at any time, with no disruption to ongoing data processing. Moreover, the RAIGE OS automatically balances all data across new nodes without interruption or impact. Within a multiple Databank system, any disk or even a Databank, can fail at any time without disruption – effectively delivering higher levels of protection than RAID 6. In addition, an innovative write-cache technology protects any in-flight data in the case of power losses during operation.

Combining Databanks together in a highly available cluster also scales performance linearly - spreading data out across multiple Databanks where it can be simultaneously accessed. The foundation of that linear scalability is a Pivot3 Multi-Path driver that distributes traffic across multiple Databanks, while also protecting against a loss of connectivity. Connectivity protection delivers fault tolerance, allowing a storage environment to sustain switch, network port, drive, or Databank failure.

Pivot3 Compared to the Criteria of Video Storage Systems

We now turn our attention to evaluating Pivot3's storage system against the key requirements that we believe a video storage solution must have to meet the new demands of digital video capture and retrieval.

√ Pivot3: Networked Storage

Pivot3 is a network attached storage system that connects to a standard gigabit Ethernet network and presents storage volumes using the iSCSI protocol. All management functions and interconnectivity between Databanks use Ethernet. The ubiquity of Ethernet throughout an enterprise allows the Pivot3 storage environment to be accessed from anywhere, while also integrating with common enterprise management and backup systems.

√ Pivot3: Easy to Use, Easy to Manage

Pivot3 optimizes operational costs as well with simplified management capabilities. In many video capture environments, operations are managed by line of business staff without specialized technology skills. Technical issues and system changes require specialized and expensive technology staff. Pivot3's easily managed interface, automated expansion, and reliance upon standard connectivity avoids the costs associated with these changes, and minimizes operational video costs just as significantly as capital costs.

Pivot3 delivers a simplified management interface that automates many aspects of administration while avoiding the presentation of complexity. Pivot3's high-levels of fault tolerance and automated cluster expansion make system configuration, failure recovery, and the expansion of available storage an easy, transparent operation. As a storage cluster is scaled across multiple Databanks, every node becomes part of the cluster allowing the entire storage system to be managed as one logical entity. Further, as a storage node scales, the system becomes completely protected from failure, allowing administrators to perform any upgrade, move, or system change without worries about affecting storage availability.

√ **Pivot3: Fault Tolerance**

Pivot3 delivers an unusual level of fault tolerance at highly competitive cost per capacity ($/TB) and price-performance ($/MB/s) metrics. By protecting data at both the node and cluster level, Pivot3 can tolerate drive or system level failures beyond what a storage array using RAID 6 can provide, while at the same time maintaining the performance and capacity advantages of traditional RAID 5 protection. Additionally, with redundancy across Databanks, Pivot3 can rebuild data from failures faster than traditional systems with RAID 5, and thus reduce the unprotected time during which a system is subject to data loss.

√ **Pivot3: Expandable Capacity**

Pivot3's RAIGE OS has been designed toward simplifying the overall configuration and expansion of the system. RAIGE enables non-disruptive capacity expansion, allowing the system to grow across added Databank nodes to form a single unified cluster. Pivot3 can support planned expansion or unforeseen future needs with the addition of capacity as it is needed. This effectively makes headaches associated with system sizing, expansion, changes in retention, or unforeseen growth a thing of the past.

√ **Pivot3: Scalable**

Pivot3 allows performance as well as capacity to be scaled non-disruptively as demand dictates. Expanding a cluster by adding a Databank increases network ports storage, processing resources, and storage capacity automatically while requiring none of the specialized skill or interruptions inherent in traditional storage. In contrast to storage systems that require purchase of larger than necessary storage controllers or subject users to complex future performance upgrades, Pivot3 grows transparently and non-disruptively to match the needs of an organization. At the same time, a Pivot3 storage infrastructure grows in availability as well, becoming increasingly resilient as storage nodes are added to a cluster.

Taneja Group Opinion

Video surveillance and capture is a burgeoning market. Many enterprises are recognizing the significant value to be had in being able to capture, review, and analyze video data, both from the perspective of increasing business capabilities and reducing liability. However, many industry sectors risk competitive disadvantage by failing to make video a strategic asset, and potentially face compliance and liability issues by poorly managing stored video data.

Enterprises that take a proactive approach to consolidating, managing, and exploring the value of video assets will reap significant competitive rewards. To that end, storage infrastructure becomes a critical decision point for capitalizing on the benefits of storing and analyzing video data. A highly available, scalable, easy to use storage system becomes a necessity in these deployments.

It is clear to us that Pivot3 has raised the bar for building video-optimized storage. Leveraging Pivot3's flexible and expandable storage design will allow enterprises to better-manage data, optimize costs, and gain more utility from their video data than competitors. If you are considering a video surveillance deployment, we strongly recommend that you consider Pivot3.