



Storage for Video Surveillance in Gaming Environments

Product Insight

John Webster
5 May 2008

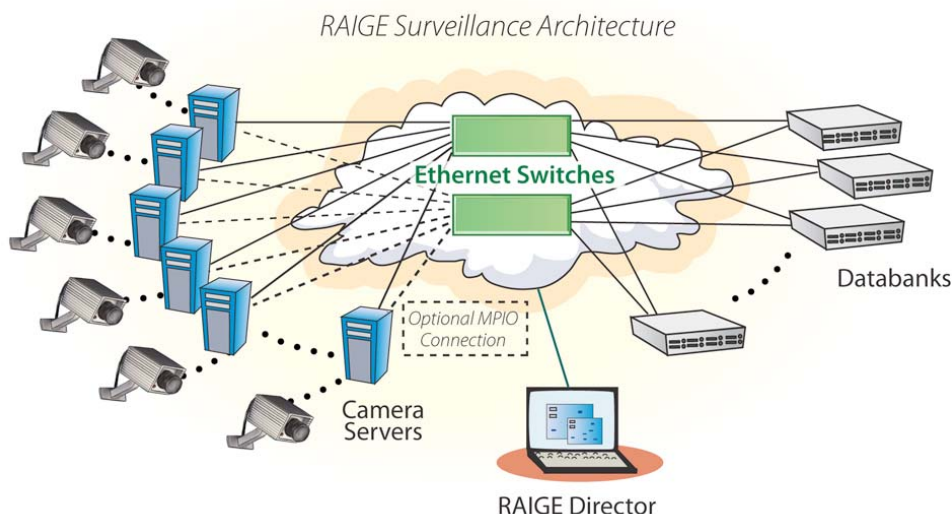
Licensed to Pivot3, Inc. for web posting. Do not reproduce. All opinions and conclusions herein represent the independent perspective of Illuminata and its analysts.

Application-specific storage will be a dominant theme during 2008. This is storage that differentiates itself from the general purpose storage—the common, high availability RAID arrays—that have traditionally prevailed within datacenters. This is storage that is optimized for a specific application or application profile. Storage for Web 2.0 is hot. Storage for video surveillance is hot as well.

Both of these application domains have at least two things in common:

1. Administrators experience phenomenal storage capacity growth rates. For example, a medium-sized casino using 100 high-definition digital video surveillance cameras can easily consume 1 petabyte of storage per month.
2. Less expensive, highly scalable storage that offers performance optimized for the required application have a distinct advantage compared to general-purpose, high-availability RAID arrays.

Here we examine the requirements for storage of high-definition images captured from IP-attached cameras positioned over key gaming tables and cash handling areas within a resort casino setting.



Capturing the Action

The video surveillance opportunity is driven by an accelerating demand for high-definition video cameras deployed over gaming tables and cash handling areas, and by state and local regulatory agencies that require image retention of one-to-two weeks. While that may not sound like a difficult set of requirements, consider that

the use of 100-1,000 cameras is not an unreasonable number for coverage of a modestly-sized casino. Now imagine ingesting and storing those video streams for up to two weeks. The need for storage array performance can render useless the storage systems that are normally supplied with these types of systems, and the capacity required can run up to the petabyte range.

Performance has always been a big selling point for upscale, datacenter RAID storage boxes. But to get both performance and availability, you generally have to pay up. Most of the casinos now appearing outside of the well-known venues (Las Vegas and Atlantic City) can't afford a petabyte of datacenter-style RAID storage dedicated solely to high-definition surveillance (let alone any other application). The answer for them is more affordable storage optimized for high-definition video surveillance.

Here, we take a closer look at storage optimized for video surveillance within the gaming industry, and specifically the experience of the Chumash Casino Resort IT staff in their deployment of a high-definition, digital video surveillance system. Their system feeds the output of IP-attached, high-definition digitized video streams into a disk storage subsystem supplied by Pivot3.

The Chumash Environment

The Chumash Casino Resort is a five-year old resort hotel and casino complex located within the Chumash tribal reserve on the central coast of California. When built, it was outfitted with surveillance cameras that are now considered obsolete. Some were analog, some were digital. None were capable of moving or zooming-in on hand movements. Nor were they capable of producing the detail required for monitoring cash windows, the counting room and vault, and gaming tables. It was determined that a number of the high-definition, IP-attached camera vendors could respond to the Chumash requirements. The storage supplied with these camera systems did not, however, meet requirements, both for performance and regulatory reasons.

High-definition digital cameras produce between four and five times the amount of data as standard cameras; each streams mega-pixel images at the rate of 15-30 per second. A storage subsystem that ingests data directly and in real time from 100 such cameras—as is the case with the Chumash implementation—requires significant front-end bandwidth.

Chumash IT administrators also had regulatory issues to contend with as they considered the upgrade. In this case, surveillance data has to be saved for 14 days. 100 high definition cameras can produce between 250 and 300 TB per week. Therefore, the minimum storage capacity required for this application was 500 TB—½ a petabyte just to “ante-up.” In addition, the camera system, once placed in operation, must remain continuously in operation.

How Pivot3 RAIGE Works for Chumash

The RAIGE architecture is composed of Databanks (iSCSI storage nodes), the RAIGE Director, and standard GigE infrastructure to connect camera servers to the RAIGE storage cluster. Cameras are connected to camera servers running video capture/management software. Pivot3's RAIGE stripes data across all physical disks in a Databank cluster. All disk drives, RAID controller functionality, and Databank nodes are virtualized by the RAIGE Operating System software to create a single pool of storage blocks.

While a RAIGE storage cluster can consist of any number of Databanks, here Chumash has installed 35, each with 12 TB of capacity. The Databanks are populated with Seagate's 7200 rpm Enterprise SATA drives. These drives were designed to accommodate the rotational vibrations, 24x7 workloads, and power management requirements inherent in the large-scale systems commonly found in surveillance applications.

Once physically established, specific Databanks in the cluster are grouped into Virtual Arrays of different capacities depending on application or user requirements. Logical Volumes of different

sizes are then created within the Virtual Arrays. Pivot3 currently supports Virtual Arrays of up to eight Databanks each, with larger configurations soon to come.¹ Architecturally, a single Virtual Array can consist of up to 32 databanks.

Scaling Performance – Performance of a RAIGE Array increases as Databanks are added and I/O operations are increasingly parallelized without “node hopping”—i.e. transferring an I/O from one node to another. Drive re-build operations (initiated after a drive failure) and data rebalancing operations (initiated when capacity is added) also see the benefit of increased parallelization as Databanks are added. Therefore a RAIGE Array can offer significant performance and scalability advantages over dual-controller arrays of comparable cost. While performance might be further enhanced by the use of host-based TCP Offload Engines (TOEs), Pivot3 does not believe they will yield any substantial benefit. In this particular application, TOEs are not used.

Scaling Capacity – Databanks can be added at anytime to both the RAIGE cluster and to any specific Virtual Array within the cluster without disruption. While Pivot3 currently supports homogeneous Databanks within a specific RAIGE Array, the RAIGE architecture supports the mixing of different databank capacities, drive types, partially populated Databanks within the cluster as well as within specific Virtual Arrays within the cluster. This gives storage administrators a significant degree of additional management flexibility over standard RAID implementations. While the Chumash installation is homogeneous today, it will likely change in the future with the addition of higher-density disks.

RAID Data Protection – RAID levels within a Virtual Array are applied and managed on a logical volume basis. All the administrative workload associated with the physical setup, configuration and management of traditional RAID sets is handled “under the covers” by the RAIGE

Operating System software. Multiple RAID types (0, 1, 5) are concurrently supported within a Virtual Array and can be dynamically modified by an administrator non-disruptively via the RAIGE Director management software (described below). The RAIGE subsystem can deliver the required performance at Chumash while set to RAID 5, eliminating the need for redundant disk.

Unlike more traditional RAID arrays that require dedicated physical spare drives, a RAIGE Array leverages the inherent parallelization of the RAIGE architecture and virtualizes spare disk capacity across all drives and Databanks within the array. RAIGE also accommodates “Databank Sparing”—a level of virtual sparing in which an entire Databank’s worth of storage capacity is reserved (distributed across all Databanks) such that if an entire Databank were to fail, enough spare capacity is guaranteed to be available so that the missing Databank’s data can be reconstructed immediately, without waiting for a replacement Databank to be brought on line.

Cluster Availability – Pivot3’s RAID-5 implementation allows a RAIGE Array to sustain both a single drive failure as well as the failure of an entire Databank within an Array. However, failure of more than one Databank automatically places that specific Virtual Array in standby mode until the Array is restored to the point where only one Databank is off-line. For applications like the one at Chumash that need maximum availability, an additional host-based multi-path I/O (MPIO) driver can be installed on each host accessing the storage cluster to provide redundant paths between application hosts and the storage cluster.

RAIGE in Software

RAIGE functionality is substantially delivered in software rather than hardware. Indeed, Pivot3 Databanks are industry-standard servers with two integrated 1 GbE ports; architecturally, up to 4 GbE ports per node are supported.

There are three software components to Pivot3 RAIGE, two required and one optional:

¹ Configurations of 2 Databanks are not supported. A single Databank RAIGE array is supported; RAID is implemented inside the Databank.

RAIGE OS – required and preloaded on every Databank prior to customer delivery. RAIGE OS delivers the core virtualization functionality of RAIGE storage clusters and RAIGE Virtual Arrays.

RAIGE Director – also required, the RAIGE Director is the administrative management interface to the RAIGE cluster and runs on a standalone workstation which can be located anywhere within the user’s network. It consists of both a graphical user interface (GUI) and a command-line interface (CLI) for automating management tasks. A single RAIGE Director instance can manage multiple RAIGE storage clusters and Virtual Arrays. Deployment of multiple instances of RAIGE Director software is also supported.

RAIGE Client Multi-Path IO – an application host-based driver that leverages multi-path I/O (MPIO) to deliver additional performance and availability. While this is an optional add-on, Pivot3 strongly recommends including this driver in RAIGE clusters for both data access (and thus application) resilience and optimal I/O bandwidth. The currently supported OS environments are Windows, Red Hat Linux, and CentOS. As mentioned earlier, multipath I/O is used in the Chumash application.

RAIGE Management

Controlling one’s RAIGE is as easy as holding your breath and counting to ten—sorry, couldn’t resist!

Seriously, Pivot3 offers a set of easy-to-use setup, configuration, performance, and on-going management screens for storage administrators through the RAIGE Director software. Management functions can also be scripted using a Command Line Interface (CLI).

Up to 128 logical volumes can be configured for each RAIGE Virtual Array. Logical volumes can be dynamically created, deleted, expanded, migrated, moved and assigned a specific RAID level using the RAIGE Director GUI. Data layout algorithms automatically stripe the data across all drives and Databanks in each Virtual Array. Data can also be non-disruptively migrated from one Virtual Array to another within the cluster as needed. Copy functions are not supported in the current RAIGE release, but are promised for a later release.

Because RAIGE presents a standard SCSI disk image to the camera servers, Virtual Arrays can be backed-up using any of the standard backup applications (Veritas NBU, Legato Networker, etc.). However, in this application, backup is not required.

Conclusion

While Pivot3 can pursue many different types of applications, the initial target market chosen for RAIGE is video surveillance. It’s a good choice, given the rapid growth in storage requirements caused by the ongoing shift from analog to digital, but also because no single storage vendor has as yet established a dominant position here.

We note that the IT administrative staff at Chumash along with their chosen systems integrator Southwest Systems, conducted an exhaustive search for storage that would adequately support their extensive deployment of high-definition digital video surveillance cameras. Pivot3 was the only vendor that could satisfy their requirements for performance and availability. As a result, Pivot3 could well establish a dominant position in surveillance storage.