

VOLUME 5 ISSUE 4

STORAGE

MANAGEMENT

SOLUTIONS™

SMS

# STORING THE HUMAN GENOME

Plus articles and case studies on:  
Data Integration Management • Electronic Vaulting  
Analytical Processing • Storage Service Providers  
Enterprise Information Portals • Storage Rights



# SAY "CHEESE" ONE TB AT A TIME

BY DAN CARSON

**F**ounded in 1995, PictureVision, now an independent wholly owned subsidiary of Eastman Kodak, is the worldwide leader in online digital image sharing and digital photo processing. Through Kodak PhotoNet online and AOL's "You've Got Pictures," consumers and professional photographers alike can make digital scans and original digital photographs available on the Internet for friends, relatives, and customers to view. PictureVision works with some 40,000 photo processing sites and all the leading independent photo labs to enable photos to be scanned, placed on CDs, or made accessible via the Web. PictureVision technology also enables customers to place online orders for reprints of photos, as well as specialty items such as prints on mugs, t-shirts, and other custom reproductions. In 1999 alone, Kodak PhotoNet, and AOL "You've

Got Pictures" accounted for over 500 million images that were either uploaded to the Internet, placed on CDs or floppy disks, or made into photo merchandise. PictureVision expects its partners to scan over one billion images in 2000, a significant portion of which will be made available for viewing on the Web.

To understand the technology issues that PictureVision faces, consider that, at any given moment, thousands of requests may be in progress for low-resolution images to be served to web browsers. Simultaneously, there may be numerous requests for high-resolution images to be downloaded for printing at a photo lab by a consumer at home. Additionally, the PictureVision network is constantly receiving newly scanned images from its photo processing partners and photo reprint orders from consumers. All PictureVision services must be available 24 hours a day, 7 days a week because it is Internet-based.

Given these parameters and the need

to store huge volumes of digital images, PictureVision requires a sophisticated, high capacity, high-speed server and storage network. Furthermore, due to continuing growth in usage, the entire computing infrastructure needs to be easily scalable and expandable.

## CHOOSING THE RIGHT SERVER AND STORAGE COMPONENTS

PictureVision works closely with Open Systems Solutions, Inc.—specialists in large-scale storage solutions—to build and evolve a server and storage infrastructure that serves current demand, yet readily expands incrementally as the popularity of online digital image management continues to grow. For servers, PictureVision uses a range of Sun Microsystems hardware. The central image repository and database is managed by a Sun Enterprise 5500. Web serving, image downloading, and online ordering applications are handled by clusters of Sun Enterprise 250 and Sun Enterprise 450 servers. Given the storage requirements of the PictureVision network, the following discussion highlights the components of the

backend storage infrastructure and some of the key factors that led PictureVision and OSSI to select and implement specific storage technology components.

### **SUPPORTING NEAR REAL TIME ACCESS**

The PictureVision network needs to store digital images in such a way that rapid access is available to service online viewing requests and requests for downloading images for printing. Optimal access times occur when a requested image is present on hard disk. However, keeping millions of existing images on disk, as well as adding millions of new images each month, is impractical and expensive, despite the incredible price/performance improvement of today's hard disk technology.

A key observation led to the chosen storage solution; the most frequently requested images are those posted on the network within the last 30 days. Therefore, images from the past month are retained on over 1TB of Sun StorEdge A3500FC RAID disk storage. The array is configured as RAID5, which stripes files across multiple disk platters to increase I/O throughput and is attached via Fibre Channel to the Enterprise 5500.

### **HIERARCHICAL STORAGE MANAGEMENT SUPPORTS LONG TERM STORAGE AND ACCESS**

After 30 days of residence on hard disk, images are migrated automatically to one of two high speed Qualstar AIT-2 tape jukeboxes. In their current configuration, the jukeboxes provide 27 TB of uncompressed file capacity with 18 drives combined accessing 480 tape cartridges. Each tape cartridge can hold 50GB of uncompressed data.


Statistically, the majority of user image requests are satisfied by images residing on the hard disk arrays. Nonetheless, a substantial number of requests do come in for images that have been migrated to tape. The AIT-2 tape format was specifically chosen because it offers the best access time for retrieving files. AIT-2 keeps its tape directory in the middle of the tape, which means that when a cartridge is accessed, the requested file is never more than half the distance of the tape. Combined with the native speed of the Qualstar drive, the jukeboxes support sub-30 second response time to the end-user for archived images.

SAM-FS software from LSC is used to manage the migration of files from hard disk to tape. The software also manages making copies of archived material to be sent offsite for disaster protection, as well as permanently

archiving the oldest material. The SAM-FS software is also responsible for the staging of requested files to the hard disk arrays on the Sun Enterprise 250 servers. This is necessary because Raster Image Processing (RIP) of image files is done "on-the-fly" on the Enterprise 250s to create the appropriate image resolution to satisfy a given user request. The ability to configure priorities for staging, archiving, and backup activities was critical in the selection of the SAM-FS software.

### **MODULAR COMPONENTS ENHANCE SCALABILITY**

The Sun StorEdge RAID disk arrays and the Qualstar AIT-2 tape libraries are designed to be easily expandable. As needed, PictureVision can simply add additional arrays or tape libraries to the existing configuration to accommodate increased storage requirements. Subsystems can be expanded without



bringing down the system, a critical factor in maintaining maximum uptime for the PictureVision network.

### **CUSTOMER SUPPORT IS IMPORTANT**

As the lead integrator, OSSI engineers were onsite during the entire implementation of the PictureVision storage components. They provided coordination and oversight of the technical support supplied by Qualstar and LSC during the installation of their hardware and software, respectively. OSSI continues to provide first-line support for the entire storage infrastructure, including on-site presence, as required. The ability for a client to call one vendor for support is a unique advantage of OSSI's storage solution expertise.

### **STORAGE WILL CONTINUE TO BE KEY FOR PICTUREVISION**

As people create more digital images, the need for space to store them will increase dramatically. Many consumers and photography professionals are becoming more comfortable with the advantages of storing these images on remote servers. Not only do they avoid a personal investment in hardware and servers, but they also gain access to new markets and new ways of sharing

**The AIT-2 tape format was specifically chosen because it offers the best access time for retrieving files. AIT-2 keeps its tape directory in the middle of the tape, which means that when a cartridge is accessed, the requested file is never more than half the distance of the tape. Combined with the native speed of the Qualstar drive, the jukeboxes support sub-30 second response time to the end-user for archived images.**

and printing images. For PictureVision, having a storage solution partner who not only understands the state-of-

the-art in storage technology, but also understands their business inside and out means they will always be able to stay one step ahead of customer demand and keep their competitors at bay.

*Dan Carson is the vice president of marketing and business development at Open Systems Solutions, Inc. (Yardley, PA).*



Reprinted by permission from the publisher of Storage Management Solutions Magazine®, Volume 5, Issue 4.

For FREE subscription information, please call 310/276-9500, ext. 36 or reply via the World Wide Web <http://www.wwpi.com>.

©2000 West World Productions, Inc.

Want the latest information?

[www.ossi.net](http://www.ossi.net)



**OSSI**  
Open Systems Solutions, Inc.

710 Floral Vale Blvd., Yardley, PA 19067  
215-579-8111 fax 215-579-8113  
1-800-898-OSSI [www.ossi.net](http://www.ossi.net)