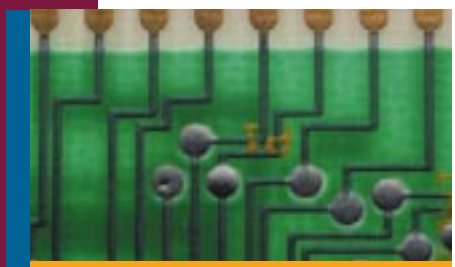

What's New in Storage?

BY STEVE PRYOR

How will some of the state-of-the-art developments in storage management affect your data center? This article explores some of the changes that have occurred recently or that are "in the pipeline."

ONE disconcerting problem with working in the data processing industry is that things change at such an alarming rate. Technology that was only recently considered state-of-the-art is now antique, along with the esoteric skills needed to operate and maintain it. In the storage arena, it sometimes seems that entirely new generations of hardware and software spring up overnight, with capabilities that make earlier systems look like they must have been put together with stone tools. Storage managers and others in data processing live constantly in an slightly awe-inspired state.

Keeping entirely up-to-date with such a rapid rate of change is impossible, of course. Yet, at least some familiarity with new developments is important, not least because if we do not take advantage of them, then someone else will. This article will explore briefly some of the changes that have occurred recently or that are "in the pipeline" and how they affect storage management in the data center.

NEW HARDWARE TECHNOLOGIES

Hardware technologies have evolved rapidly in order to meet three of the most important needs of the data processing environment: high performance, high capacity and high availability. The past few issues of the "Storage Strategies" column have dealt with minimizing application downtime during backup or copy operations through the use of hardware capabilities such as EMC Timefinder, IBM/STK SnapShot Copy and HDS Data-Plex. Other hardware technologies address the pressing issue of accessing data across multiple platforms, such as between OS/390 and UNIX or Windows NT. The IBM InfoSpeed data gateway device, for example, attaches to an S/390 server via ESCON chan-nels on one side and to open systems platforms such as Sun or HP via SCSI channels on the other side. The InfoSpeed Express Data Mover (PDM) software runs on both the host S/390 system and on the open systems platforms to move data between the desired systems. Such movement might be done to "clean up" data when moving it to a different operating system, for example.

Two other hardware devices from IBM and similar devices from other manufacturers also deal with the rapidly growing open systems storage environment. The IBM 9399 Cross Platform Extension (XPE) allows UNIX and NT data to reside on RAMAC array devices such as the RAMAC Virtual Array (RVA). The open systems data can be backed up or restored (at the 3390 device level) with the more mature storage management facilities of the mainframe. Frozen point-in-time copies of the open systems data can be made via SnapShot. The IBM Versatile Storage Server is an open systems storage device that is intended to centralize and share disk storage among many different types of open systems servers.

Tape devices are also an important and rapidly-advancing component of data center operations. The Virtual Tape Server (VTS) provides a means for high-speed access to tape data via an internal DASD cache. In addition to largely eliminating tape mount time, the VTS stacks multiple "virtual" tape volumes on a single physical tape, fully utilizing the enormous capacity of the 3590 Magstar tape drive.

All of these new technologies introduce change — and that means change in data center storage management methods. Disaster recovery preparation must include ensuring

the availability of new types of hardware. Rules must be established for the expiration of backups of open systems data that might now reside on the mainframe. New concepts such as logical tape volumes must be understood. VTS resources such as the Tape Volume Cache must be managed to maximize cache "hits" and prevent performance-degrading overflow and so on.

SOFTWARE

Using and managing storage hardware requires software. Like hardware, software has also evolved to meet the needs of the data center for performance, capacity, and reliability. On September 29, IBM announced that one of the components of the DFSMS Optimizer, the DFSMShsm Monitor/Tuner, is being "unbundled" from the Optimizer and will be available as a separate product named the StorWatch DFSMShsm Monitor. The StorWatch DFSMShsm Monitor allows users to view the status of HSM tasks on all systems, set warning thresholds for potential error conditions, and take other actions such as starting or stopping DFSMShsm activity via a TCP/IP connection with a Windows or OS/2 workstation. Of course, IBM is not the only vendor introducing changes in the way storage is managed. Amdahl's Transparent Data Migration Facility (TDMF) creates point-in-time duplicates of application data similar to those created by hardware facilities, but without requiring any particular manufacturer's type of device. A "master" task and a set of "slave" tasks, which may be on different OS/390 images, communicate to set up "migration pairs" of source and target DASD volumes. Data is copied from the source to target volumes while monitoring I/O to the source volume. If a change to the source data is detected before all of the data has been sent to the target volume, then the changed tracks are re-copied. Eventually, all of the data is copied and no further refreshes are needed. At this point, the source and target volumes are duplicates of one another. The target volume can be used to provide an independent copy of the data for testing or backup, allowing operations to continue uninterrupted on the production volumes. The target volume can also dynamically take over access to the application data if the object was to migrate to new devices.

Naturally, TDMF comes with its own set of complexities. Communication datasets (COMMDSs) must be set up to allow communication between the master and slave tasks. Master and slave JCL must be created and maintained. TDMF sessions may need to be monitored for activity or to change volume synchronization intervals. TDF documentation is available at www.amdahl.com/support/datacenter/tdmf/pub.

In the storage arena, it sometimes seems that entirely new generations of hardware and software spring up overnight, with capabilities that make earlier systems look like they must have been put together with stone tools.

Changes are also introduced to the data center as operating system and program product levels advance. New OS/390 releases now occur on a regular basis every six months, and each contains important new features waiting to be exploited. DFSMS version 1.4 introduces a slew of new items pertinent to storage managers. These include enhancements to VSAM that raise the limit on the number of extents that a component may occupy to 255, and new data class parameters (RECORD ACCESS BIAS) and JCL parameters (AMP=ACCBIAS=…) that allow the system to determine the optimal buffering for VSAM datasets. Improvements to DFSMShsm include (finally) allowing the creation of duplex copies of tapes at back up time, stacking of ABARS backup tapes, and the use of the OPTIMIZE keyword for ABARS backups. Looking only slightly further down the road, DFSMS version 1.5 will include enhancements that take advantage of sysplex facilities to allow such features as DFSMShsm failover protection — promotion of a DFSMShsm secondary host to primary if the DFSMShsm primary host fails. S/390 UNIX Services (OpenEdition) HFS datasets will no longer be limited to a single DASD volume. In short, there are a lot of changes to keep up with as hardware and software technologies progress.

NETWORK STORAGE SYSTEMS

With the explosive growth in the number and kinds of storage devices available, many of which are attached to corporate networks but are not under central control, just figuring out how much storage is accessible and how much is used (and when it will run out) can be a difficult problem to solve. A basket of incompatible vendor-specific tools for managing diverse types of hardware only adds to the complexity of the situation.

In an attempt to bring order to some of the chaos introduced by all this change, a new discipline known as Enterprise Storage Resource Management (ESRM) is beginning to evolve. ESRM marries the facilities of networks with those of storage management software to provide a way of monitoring and managing the state of storage resources across the entire corporation. ESRM software includes management and reporting components that bring together information from all of the storage devices attached to the network. The information is provided by software "agents" that run on the various systems attached to the network and present consistent information back to the management component. Armed with this information, the status of widely dispersed, diverse storage resources can be watched, adjusted and subjected to automation, all with the goal of enforcing high-level, enterprise-wide storage policies.

IBM's entry in the ESRM area is StorWatch, which includes agents to perform services such as discovering the location of servers and storage devices and issuing alerts when event thresholds are encountered. Other components, such as the StorWatch Versatile Storage Specialist and the StorWatch Serial Storage Expert provide configuration management services for compatible devices. The StorWatch Reporter component uses a web browser interface to provide access to capacity and utilization information. Other agents and administration functions already exist or are planned.

All indications are that we can expect the current trends in the storage environment to continue. More and more storage will be needed, larger and faster devices will continue to appear, and more software, including network-enabled tools, will be built to manage them. And for those of us looking for a breather, we're probably out of luck — the rate of change will continue to be alarming. ts

NaSPA member Steve Pryor has more than 15 years of experience in storage management, disaster recovery, software development, and technical support. Steve can be contacted via the Internet at pryor@atlanta.com.

