How Fast is Fast?
Point-in-Time Backup Technology

BY STEVE PRYOR

Backups are important — they are the storage administrator's first line of defense against lost or corrupted data. For most storage administrators, ensuring that the backups run successfully every night, and seeing that a copy of the backed-up data is transported safely offsite, is one of the highest-priority tasks. The availability of reliable backups can mean the difference between success and failure of the business when disaster strikes.

Yet, backing up data is essentially non-productive. It doesn't add any value to the data — it merely makes a copy to act as an insurance policy against problems. Worse yet, backups are time-consuming and disruptive to applications, which must be quiesced so that the datasets included in the backup are logically consistent.

MORE DATA TO BACK UP

The capacity of DASD storage devices has grown exponentially over the last few years, and there is now much more data to back up. Year 2000 testing, data mining databases, and the growth of Internet applications have all contributed to the need for larger backups, which take longer to complete. With backups taking ever longer, the requirement to quiesce applications for the duration of the backup becomes intolerable. This is particularly true as business becomes more globalized, and the window during which applications can be halted becomes ever smaller.

FAST BACKUP TECHNOLOGIES

A number of hardware technologies have arisen which allow the creation of near-instantaneous copies of large amounts of data. These may be physical copies, as in the case of 3990 dual-copy volumes, or they may be logical copies such as those provided by the SnapShot facility of the Ramac Virtual Array (RVA) from IBM. In either case, software technologies have also been developed to back up the duplicates created by the hardware. Since the duplicate is (in most cases) independent of the original data, the impact on applications is much less — they need be quiesced only for the few minutes or seconds that it takes the hardware to make the duplicate, rather than the hours that are required for the physical backup.

Concurrent Copy

One of the first hardware technologies to allow for the creation of "fast" backups was the Concurrent Copy facility of the IBM 3990-6 controller. To create a backup using this feature, the CONCURRENT option is specified on the DFSMSdss DUMP or COPY command, and DFSMSdss invokes a utility called the System Data Mover (SDM) to establish a Concurrent Copy "session." SDM then makes a record (in a "side file" in the 3990 cache) of which tracks are occupied by datasets that are to be backed up. At this point, even though no data has actually been moved, the backup is logically complete, and applications may be restarted. If any of the tracks involved in the Concurrent Copy session are updated, they are first copied to the 3990 controller cache, so that the backup process still has access to the original, unchanged data. Should the 3990 cache fill up due to a large number of updates, MVS data spaces are used to hold the tracks.

Concurrent Copy has the advantage of using hardware that is commonly available and is useful for applications such as DFSMShsm control dataset backup, where critical program functions are disabled until the backup is complete. There are several drawbacks to the use of Concurrent Copy, however. Since the disk controller cache is being used for the backup, rather than its originally intended purpose of caching, performance may be adversely affected. Concurrent copies should normally be performed only when system activity is low, which is at odds with the objective of limiting application downtime. More importantly, if an error occurs during the physical (data movement) portion of the backup, the Concurrent Copy session is lost. It's impossible to restart the backup from the same point in time, since Concurrent Copy does not duplicate the original data prior to beginning the backup.

Other technologies that create or maintain duplicate copies of data provide other avenues for minimizing the impact on applications caused by the need to make backups.

The SnapShot Facility of the IBM RVA

The optional SnapShot facility of the IBM Ramac Virtual Array (RVA, nee the StorageTek Iceberg) can create instantaneous duplicates of volumes or datasets by duplicating pointers maintained in its internal tables. The duplicates can be backed up independently of the original data. Until very recently, the System Data Mover did not support SnapShot, and installations that used Concurrent Copy were unable to perform point-in-time backups on data that had been migrated to RVA devices. For those who had come to depend upon Concurrent Copy for such things as DFSMShsm CDS backup, this was a real problem, and APAR OW30915 was released to describe how CDS backup could be done with SnapShot rather than the CDSVERSIONBACKUP command. Other migration issues also arise with the RVA, one of the most annoying being the issuance of DFSMSdss error messages (ADR379E or ADR299E) due to the fact

© 1998 Technical Enterprises, Inc. Reproduction of this document without permission is prohibited.

that the RVA does not have any alternate tracks. When a full-volume copy or restore was performed to move data from 3390 volumes to the RVA, the format-4 DSCB incorrectly indicated the existence of alternate tracks, and it was necessary to do an ICKDSF REFORMAT REFVTOC to correct the problem before anything else could be done to the volume. A recent DFSMSdss APAR (PN89166) and others for DFSMS addressed this problem.

The EMC Symmetrix

The EMC Symmetrix also includes a facility for making duplicate backup copies of DASD volumes. These duplicates are called Business Continuance Volumes (BCVs). Unlike the RAID-1 duplicates that can also exist in the Symmetrix to protect against hardware errors, the BCVs have their own device addresses and can be linked to or detached from ordinary volumes via the EMCTF utility. The ESTABLISH command causes EMCTF to copy data from the original volume to the BCV, and once all the data is copied, updates to the original volume are mirrored to the BCV so that it remains in sync with the original. The EMCTF SPLIT command can then be used to detach the BCV, which can be backed up independently. A similar feature is available for the Hitachi Data Systems HDS 7700, which uses the HDS HRC Remote Copy utility.

SPECIAL CONSIDERATIONS FOR BACKING UP DUPLICATE COPIES OF DATA

The backup of duplicate copies of data has some special considerations. Except for dataset-level SnapShot, the duplicate is a copy of an entire volume. Usually, the duplicate has the same volume serial number as the original volume, which remains online so that the applications which use it can continue to run. In order to back up the duplicate, therefore, it must either be brought online to another system and backed up from there, or the volume must be relabeled with ICKDSF before being brought online to be backed up. If it is relabeled, however, the VVDS name will no longer match the volume serial number, and neither the VSAM datasets nor the SMS-managed datasets will be accessible. This is also true for SnapShot volume copies where COPYVOLID(NO) is specified, resulting in a target volume whose volser differs from the original. An alternative is to leave the duplicate volume offline and use Innovation Data Processing's FDR InstantBackup, which can back up data on offline volumes.

The area of "fast" backup is still new and evolving. Problems occasionally arise, as in the case of APAR OW26683, which fixed a problem involving data corruption when PDSEs were duplicated by SnapShot. Others will no doubt arise, and the capabilities of products that perform such minimally disruptive backups will continue to improve. For the storage administrator who is facing the dual pressures of the need for backup and the need to keep applications available, an understanding of the hardware and software technologies involved can be invaluable. 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.

©1998 Technical Enterprises,Inc. For reprints of this document contact sales@naspa.net.

TECHNICAL SUPPORT
Supporting Enterprise Networks and Operating Environments
SUPPORT