"Instant-Copy" Backups

BY STEVE PRYOR

For the past couple of months, I've discussed the new "instantaneous copy" capabilities that are available with recent advances in DASD hardware, such as the SnapShot copy capability of the IBM Ramac Virtual Array and the Timefinder function of the EMC Symmetrix. As we've seen, these technologies give the storage administrator new opportunities to carry out his or her duties, with much less of an impact on the "real" production work of the data center. One of the most important duties is ensuring that daily backups run successfully within the allotted time so that critical data can be recovered in the event of an error. This month, I'll conclude this series with a discussion of how the major DASD management products have provided support for these point-in-time copy functions.

DFSMSdss support for the IBM RVA/STK Iceberg SnapShot feature requires DFSMS version 1.3 and APARS OW29881, OW29883, and several others. This support allows DFSMSdss to invoke SnapShot copy for both the DUMP and COPY commands. For COPY, nothing special needs to be done; if the source and target of the copy both happen to reside in the same RVA, then SnapShot is automatically invoked by DFSMSdss to copy the tracks directly from one virtual volume to another. (Actually, of course, only a copy of the pointers to the tracks is made in the RVA's Functional Track Table for the target volume.) For DUMP, SnapShot support is implemented via almost exactly the same mechanism that was supported for the original 3990 Concurrent Copy hardware, that is, via the CONCURRENT keyword. If Snapshot is invoked through the use of the CONCURRENT keyword, the resulting dump is taken using "Virtual Concurrent Copy." Instead of setting aside space within the 3990 cache controller, the Virtual Concurrent Copy function invokes SnapShot to copy the source data into one or more Working Set Data Sets (WSDSs). The WSDSs serve the same function as the cache sidefile did in the 3990 – they act as a temporary holding area to allow the backup to logically complete in a few seconds, even though the backup may not physically complete for some hours.

The appearance of new hardware technologies that allow near-zero application downtime is an advance that allows installations to move still closer to the holy grail of 24x7 continuous availability.

The storage administrator is responsible for creating and cataloging the WSDSs, which must exist before DFSMSdss can use Virtual Concurrent Copy. These datasets are ordinary sequential datasets, which may or may not be SMS-managed, and are named SYS1.ANTMAIN.Ssysname.SNAPnnnn, where 'nnnn' is an ascending numeric value. DFSMSdss will use as many or as few WSDSs as are needed to contain all of the tracks being dumped. At a minimum, one WSDS will be needed per device type per RVA, and each DFSMSdss job that runs on a system will need its own set of WSDSs. One way to ensure that there is sufficient WSDS space available is to allocate a full-volume-size WSDS for each virtual volume in the RVA. Since these datasets seldom occupy any real back-end storage space, they do not seriously impact Net Capacity Load (NCL), although they do take up virtual space in the RVA.

Using the existing Concurrent Copy architecture has the advantage of not requiring any changes for installations that have been able to take advantage of the Concurrent Copy feature of the 3990 controller. If datasets currently being backed up via Concurrent Copy are moved to RVA devices, then Virtual Concurrent Copy (SnapShot) can be used without any further effort on the storage administrator's part. Existing console automation, which relies on the ADR734I message to indicate the completion of a logical dump, for example, can remain unchanged. However, this also means that one of the significant disadvantages of Concurrent Copy – the inability to "freeze" a point in time – is continued. Even though SnapShot will make a complete copy of data being backed up, that data may be spread across several WSDSs and is not accessible outside of DFSMSdss. Thus, the point in time is lost – the backup cannot be restarted with the same data – if the physical movement of data to tape fails.

CONSIDERATIONS FOR IMPLEMENTING SNAPSHOT

There are a few other considerations the storage administrator should take into account when implementing SnapShot. One is the introduction of new fields in the SMS storage class that influence the selection of volumes on which a dataset can reside. The BACKUP and VERSIONING parameters of the ACCESSIBILITY attribute can be used to determine whether a new dataset is allocated on a SnapShot-capable device. Another consideration is whether DFSMSdss must invoke a utility such as IDCAMS in order to back up a

dataset, in which case SnapShot cannot be used. Also, SnapShot does not support VSAM datasets that are defined with the IMBED and REPLICATE parameters; these types of datasets are copied using conventional techniques.

Installations that use Innovation Data Processing's FDR InstantBackup can also take advantage of the instantaneous-copy facilities provided by the RVA and other such devices. When copying datasets from one disk to another using FDRCOPY, SnapShot will be invoked if the source and target dataset reside in the same RVA. The processing used by the FDR system for dumping data is somewhat different from that used by DFSMSdss in that no WSDSs are used. Instead, an offline copy of an entire volume is made using program FDRSNAP (or IXFP). An ordinary FDR or FDRDSF backup can then be run, specifying the original online volume serial number, with the SNAP=USE or SNAP=(USE,REL) parameter. Even though the original source volume serial number was specified for the backup, the actual backup will be of the offline, snapped copy. Since the snapped copy reflects the original volume, rather than a collection of tracks as in a WSDS, the backup can be restarted if the physical copy to tape should fail. The REL operand of the SNAP parameter allows Deleted Data Space Release (DDSR) to be run after the physical backup has completed, in order to reduce NCL if many tracks have been updated between the logical and physical completion times.

The appearance of new hardware technologies that allow near-zero application downtime is an advance that allows installations to move still closer to the holy grail of 24x7 continuous availability.

FDR InstantBackup also provides a similar offline backup of the duplicate point-in-time copies of EMC Symmetrix Business Continuance Volumes (BCVs) and Hitachi Data Systems HDS 7700 duplicates, and can use Snapshot or the internal tracks-copy feature of the Symmetrix to move data during a COMPAKTOR volume reorganization. Although there is no longer any need to reorganize volumes for performance reasons, reorganizations may still be useful to reduce the number of extents occupied by datasets and provide enough contiguous (virtual) free space on volumes to prevent out-of-space errors.

The appearance of new hardware technologies that allow near-zero application downtime is an advance that allows installations to move still closer to the holy grail of 24x7 continuous availability. For the storage administrator, ensuring that the backup software takes full advantage of these powerful new capabilities is an important part of this effort. 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.

