# Picturing Virtual DASD — SnapShot Copy

BY STEVE PRYOR

Last month's column discussed the availability of new hardware and software that makes point-in-time backup a possibility. This month, I'd like to focus more specifically on one of these platforms, the popular IBM RAMAC Virtual Array, or RVA, and its impact on the storage administrator.

The RVA is one of several types of RAID storage devices developed by StorageTek Corporation and now marketed by IBM. In addition to the RVA (originally the Iceberg), the Kodiak device is now sold as the RAMAC Scalable Array, and the Artic Fox has become the RAMAC Electronic Array.

Each of these devices provides the large capacity and redundancy that are the hallmark of RAID technology. The RVA, however, includes a unique feature that is of particular importance to the storage administrator, namely the SnapShot copy feature, which allows the almost instantaneous duplication of anything from a single dataset to the entire contents of an RVA subsystem.

This ability to "freeze" large volumes of data at a point in time without any serious impact on applications makes the RVA very useful for such tasks as Year 2000 testing, data mining applications, database duplication, and other tasks that require additional copies of production data. IBM estimates that about 70 percent of the RVA's currently being shipped include the SnapShot feature.

SnapShot takes advantage of the Log Structured File architecture that is used to store data on the RVA. Although the DASD volumes in the RVA appear to MVS to be ordinary 3380 or 3390 volumes, these are only "functional" devices, represented by a table (the Functional Device Table) within the RVA. The tracks on each virtual device are represented by another table (the Functional Track Table). These tables are then mapped to the actual devices and sectors of the real IBM Ultrastar 2XP disk volumes that make up the array.

## FEATURES OF THE RVA

Under the Log Structured File architecture, the back-end storage on the RVA is used in a linear fashion; that is, tracks on the RVA are never rewritten. Instead, any information written to the RVA, whether to create new data or to update existing data, is always written to a new location on the real devices. (From the MVS point of view, of course, it appears that an existing 3380 or 3390 track was written or updated). The use of the Log Structured File architecture provides several advantages. First, tracks on the apparent MVS volumes that are allocated but unused do not take up any space on the real devices that make up the back-end storage. The same is true for tracks that are unallocated on the MVS volumes. While entries exist in the Functional Device Table and Functional Track Table for these tracks, they are not mapped to actual storage until data is placed on the virtual track. Secondly, all data actually written to the back-end storage is compressed. These features allow the RVA to hold more virtual data than the back-end storage actually has the capacity for.

The percentage of the RVA that actually holds real data, as compared to the total capacity of the unit, is known as Net Capacity Load (NCL), and is an important measure used in managing RVA storage. Equally important, however, is the amount of "collected" free space. Because new or updated data is always written to a new location on the real devices that make up the back-end storage, a mechanism is necessary to free up the old locations when new sectors are written. Physical space in the RVA can only be used in units of "array cylinders" (which do not necessarily correspond to real device cylinders), so a free-space collection task runs periodically within the RVA to gather these unused sectors into array cylinders. Maintaining enough free array cylinders is important because the less free space available, the more time the RVA must devote to free space collection in order to ensure that new or updated virtual tracks can be written. IBM's rule-of-thumb recommendations for the RVA are to maintain the NCL at 85 percent or less and the collected free space at 10 percent or more. If there is a lot of write activity to the device, then a lower NCL and more collected free space may be needed.

In addition to the free space collection task that runs within the RVA, the storage administrator may occasionally need to run an MVS task to perform Deleted Data Space Release (DDSR). This will notify the RVA that space on the virtual devices is no longer in use and may be reclaimed. Dynamic DDSR runs automatically each time a dataset is deleted or has its unused (virtual) tracks released. Interval DDSR runs as a batch job under IXFP, the IBM Extended Facilities Product. Since Interval DDSR holds a VTOC enqueue while it identifies the free virtual tracks, it should not be run indiscriminately against a large number of volumes. Interval DDSR may be needed, however, in situations where dynamic DDSR has not been invoked, such as after a full-volume restore or after establishing an RVA as the target of an Extended Remote Copy (XRC) pair.

IXFP is the main vehicle used by the storage administrator to communicate with the RVA. Besides performing Deleted Data Space Release, IXFP provides reports on the NCL, amount of collected free space

and other RVA statistics. Another important function of IXFP is to create "instant" copies of data using the SnapShot feature that is made possible by the RVA's Log Structured File architecture. When IXFP is invoked via a batch job, REXX EXEC, or TSO command, the SNAP DATASET or SNAP VOLUME commands instruct the RVA to make a duplicate copy of the pointers in the Functional Track Tables. Once this is done, there are effectively two copies of the data, since there are two sets of pointers to the same physical location in the array. If either copy of the data is updated, the data is written, as always, to a new location in the array, and the appropriate Functional Track Table is updated.

Since only pointers, rather than real data, are copied, a SnapShot copy is extremely fast. It may take only a few seconds to make a copy of an entire virtual 3390 volume. The duplicate copy does not take up any physical space in the array. In fact, even though an entire DASD volume may have been duplicated, the Net Capacity Load on the RVA does not change, since there is no more real data on the back-end devices than there was before the copy. The NCL only begins to increase when the original data (or the copy) is changed.

## CONSIDERATIONS

The storage administrator must consider several factors when implementing SnapShot. First, the target of the copy must reside in the same RVA unit as the source. This means that SMS storage groups may need to be structured so that volumes which may be involved in a SnapShot operation are assigned to the same storage groups. When IXFP is installed, a "volume preferencing" function is also inserted into DADSM, which allows IXFP to attempt to select devices within the same RVA when allocating a new dataset as a SnapShot target. If no volume in the same RVA is eligible to hold the target dataset, IXFP will either fail the job or invoke a data mover (typically DFSMSdss) to copy the data rather than performing a SnapShot copy.

If the target of a SnapShot copy already exists prior to the copy operation, then its attributes, especially non-VSAM DSORG or VSAM CI size, etc., must match that of the source dataset. Some VSAM attributes, such as IMBED and REPLICATE, are not supported for dataset Snapshots (although such datasets may be included in a volume Snapshot, of course). Although they are supported on the RVA, extended-format VSAM datasets and those with alternate indexes are not currently eligible for Snapshot copy, though this may change very shortly.

For the storage administrator, the most important aspect of the SnapShot feature is that creating a copy of large amounts of data no longer requires large amounts of application downtime in order to get a logically consistent copy. Only slightly less important is the fact that the additional copy does not take up any additional real disk space unless the application has a high rate of write activity. Given the ever-increasing amount of storage that has to be managed, an understanding of RVA technology and the SnapShot copy function is critical for the storage administrator. 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.**