Technical Brief

Working with the Two Terabyte Volume Limit in Some Operating Systems

Many current operating systems have a 2 TB addressing limit. ATA drive capacity continually increases, and so does the overall capacity of the ATTO Diamond Storage Array.

Drive Size	Diamond Capacity
80 GB	1.92 TB
120 GB	2.88 TB
160 GB	3.84 TB
180 GB	4.32 TB
250 GB	6 TB
300 GB	7.2 TB
500 GB	12 TB

At this time, most host operating systems can only address up to 2 TB per LUN. This is becoming a problem for users with large capacity storage arrays, such as the ATTO Diamond Storage Array. This Technical Brief will discuss different ways for working around this operating system limit with the Diamond Array.

Dividing the Diamond Array into multiple physical LUNs

The simplest method to accommodate the 2 TB limit is to divide the Diamond Array's 24 physical drives up into separate Virtual Disks using the ATTO QuickRAID™ command. For example, a 6 TB Diamond Array can be broken up into three separately addressable RAID 0 Virtual Disks, or LUNs, by issuing a *set QuickRAID0 3* command. Each Virtual Disk would consist of eight drives with a capacity of 2 TB. If RAID 5 is preferred, simply issue a *set QuickRAID5 3* command to create three separate RAID 5 LUNs. Again, each LUN would consist of eight drives, but only six would be addressable. The other two are for parity. This gives a capacity of 1.5 TB per LUN.

The options available for RAID 0 configurations include 1, 2, 3, 4, 6 or 12 Virtual Disks. To determine the minimum number of Virtual Disks required to compensate for the 2 TB limit, simply divide the total capacity of the Diamond by two and round up to the next highest whole number.

Diamond Capacity	Minimum # of RAID 0 Virtual Disks
1.92 TB	1
2.88 TB	2
3.84 TB	2
4.32 TB	3
6 TB	3
7.2 TB	4
12 TB	6

If a RAID 1 configuration is required, there is no issue unless the Diamond Array has drives that are over 1 TB in size. For RAID 1, the Diamond Array concatenates both drives on a sled and then mirrors them with both drives on the adjacent sled. The addressable capacity of each RAID 1 Virtual Disk is the capacity of two drives. Issuing the *set QuickRAID 1* command will result in the creation of six individual Virtual Disks, with four drives each.

RAID 10 is similar to RAID 1 in that half of the Diamond Array's capacity is addressable while the other half is used as a copy. For RAID 10, the Diamond Array takes the previously described RAID 1 LUNs and stripes them together. The options available for a RAID 10 configuration include 1, 2 or 3 Virtual Disks. The capacity for each of these LUNs can be calculated by taking the total capacity of the array, dividing it in half, and then dividing that by the number of configured Virtual Disks. Be careful to ensure this capacity stays under 2 TB.

Diamond Capacity	Minimum # of RAID 10 Virtual Disks
1.92 TB	1
2.88 TB	1
3.84 TB	1
4.32 TB	2
6 TB	2
7.2 TB	2
12 TB	3

Calculating the minimum number of RAID 5 Virtual Disks to be sure the capacity of each LUN remains under 2 TB is a little more complicated because of the parity drives. The Diamond array will allow the configuration of 1, 2, 3 or 4 RAID 5 Virtual Disks. Each of these configurations will have two parity

Diamond Capacity	Minimum # of RAID 5 Virtual Disks
1.92 TB	1
2.88 TB	2
3.84 TB	2
4.32 TB	2
6 TB	3
7.2 TB	3
12 TB	4

drives per LUN, the rest are addressable. For example, a three LUN RAID 5 configuration will divide the 24 drives up into three Virtual Disks of eight drives each -- six addressable and two parity. The table on the previous page can be used to determine the minimum number of RAID 5 LUNs to ensure the 2 TB limit is not exceeded.

RAID 5 users may not wish to divide the Diamond Array into smaller Virtual Disks because overall capacity is lost as the amount of parity space increases. With each LUN requiring two drives for parity, up to eight drives worth of capacity could be lost. To compensate, the Diamond Array also has a Virtual Disk Partitioning feature.

Disk Partitioning

Partitioning offers users the ability to divide the array storage space into multiple addressable sections, serving as logical units. For example, the Diamond Array can be configured for one RAID 5 Virtual Disk. This would contain 22 addressable drives and two parity drives (which is the most capacity-efficient RAID 5 configuration). This Virtual Disk can then be split into anything up to 16 equally sized Partitions. Each partition will be addressable as a LUN.

The figure below depicts a four drive Virtual Disk that was sliced into three equally sized contiguous partitions. If each drive was 640 GB in size, the Virtual Disk would be 2.56 TB. With partitioning, each LUN is now 853 GB (2.56 TB/3), well within the 2 TB addressable limit.

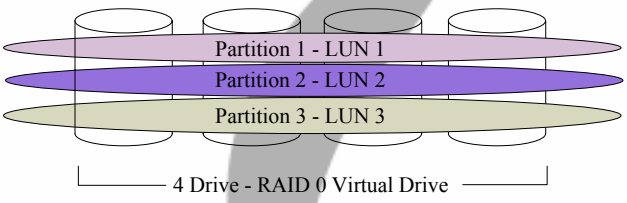

As another example, consider the RAID 5 case in which 250 GB drives are used. This 6 TB array would have to be broken into three Virtual Disks to stay under the 2 TB addressing limit.

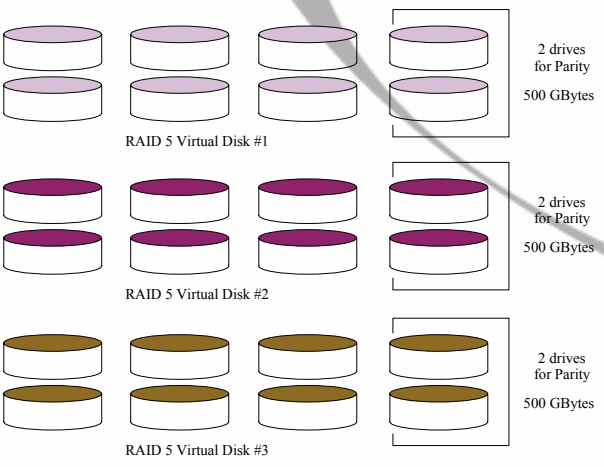

Three Virtual Disks require six drives worth of capacity for parity. The loss of 1.5 TB of capacity may be a concern for some users.

An alternative would be to create only one RAID 5 Virtual Disk and then create three partitions. This results in only two drives, or 500 GB of unavailable capacity.

The trade off is performance. With multiple partitions being distributed across all of the drives, the drive head must move back and forth across the platter if multiple partitions are being accessed simultaneously. The additional seek time will result in lower MB and I/O per second performance.

Summary

The QuickRaid facility allows the storage to be divided into more, smaller, physical LUNs. In Raid 5 this results in a loss of user data capacity. Partitioning can increase storage efficiency by providing the option of creating more LUNs without requiring the creation of lower capacity RAID groups. The trade off here is if multiple partitions on the same virtual RAID set are being accessed at the same time, disk seek time can affect performance as the drive head moves back and forth between the multiple partitions.

Each user needs to understand his/her requirements for capacity versus performance when determining which configuration to choose. Most users would most likely benefit from a combination of the two methods for managing the two Terabyte limit.

Refer to the Diamond Storage Array Installation, Operation and Maintenance Manual for further details on the RAID configurations. The *DriveInfo, VirtualDriveInfo* and *PartitionInfo* CLI commands can be used to obtain individual drive capacity, RAID configuration, partitioning and LUN relationship information.

Corporate Info

ATTO Technology, Inc., is a global leader in Fibre Channel, IP and SCSI storage and storage-infrastructure solutions for Enterprise and Workgroup environments. Since its founding in 1988, ATTO has been designing, manufacturing and marketing award-winning solutions specifically for data-intensive applications. ATTO distributes its products worldwide through original equipment manufacturers (OEMs), distributors, VARs and System Integrators.

ATTO Technology, Inc.

Corporate Headquarters
155 CrossPoint Parkway
Amherst NY 14068

phone: 716.691-1999
fax: 716.691-9353
www.attotech.com.

©2003 ATTO Technology, Inc