

POWERSVAULT MD3000 AND MD3000i

ARRAY TUNING
BEST PRACTICES

dell.com/PowerVault



Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

DISCLAIMER

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

For more information, contact Dell™.

Information in this document is subject to change without notice.

<http://www.dell.com>

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

Table of Contents

1	AUDIENCE AND SCOPE	4
2	PERFORMANCE TUNING OVERVIEW	4
2.1	COMPONENTS THAT INFLUENCE STORAGE PERFORMANCE	4
2.2	BASIC APPROACH TO PERFORMANCE TUNING	4
3	APPLICATION SOFTWARE CONSIDERATIONS.....	5
4	CONFIGURING THE MD3000/MD3000I	6
4.1	DETERMINING THE BEST RAID LEVEL	6
4.1.1	<i>Selecting a RAID Level - High Write Mix Scenario.....</i>	<i>8</i>
4.1.2	<i>Selecting a RAID Level - Low Write Mix Scenario.....</i>	<i>8</i>
4.2	CHOOSING THE NUMBER OF DRIVES IN A DISK GROUP.....	8
4.3	VIRTUAL DISK LOCATION AND CAPACITY	9
4.4	VIRTUAL DISK OWNERSHIP.....	10
4.5	CALCULATING OPTIMAL SEGMENT AND STRIPE SIZE	10
4.6	CACHE SETTINGS.....	12
4.6.1	<i>Setting the Virtual Disk-Specific Write Cache and Write Cache Mirroring.....</i>	<i>12</i>
4.6.2	<i>Setting the Virtual Disk-Specific Read Cache Pre-fetch.....</i>	<i>12</i>
4.6.3	<i>Setting the Storage Array Cache Block Size</i>	<i>13</i>
4.7	TUNING USING ARRAY PERFORMANCE DATA.....	14
4.7.1	<i>Collecting Performance Statistics</i>	<i>14</i>
4.7.2	<i>RAID Level</i>	<i>14</i>
4.7.3	<i>I/O Distribution</i>	<i>15</i>
4.7.4	<i>Stripe Size.....</i>	<i>16</i>
4.7.5	<i>Write Algorithm Data</i>	<i>19</i>
4.8	USING THE CLI PERFORMANCE MONITOR.....	21
4.9	OTHER ARRAY CONSIDERATIONS	22
4.9.1	<i>Global Media Scan Rate.....</i>	<i>22</i>
4.9.2	<i>Setting the Virtual Disk-Specific Media Scan</i>	<i>22</i>
4.10	PREMIUM FEATURE PERFORMANCE	22
4.10.1	<i>Getting Optimal Performance from Snapshot.....</i>	<i>22</i>
4.10.2	<i>Getting Optimal Performance from Virtual Disk Copy.....</i>	<i>22</i>
5	CONSIDERING THE HOST SERVER(S).....	23
5.1	HOST HARDWARE PLATFORM	23
5.1.1	<i>Considering the Server Hardware Architecture.....</i>	<i>23</i>
5.1.2	<i>Sharing Bandwidth on the Dell™ MD3000i with Multiple NICs</i>	<i>23</i>
5.1.3	<i>Sharing Bandwidth with Multiple SAS HBAs.....</i>	<i>24</i>
5.2	CONSIDERING THE SYSTEM SOFTWARE	25
5.2.1	<i>Buffering the I/O.....</i>	<i>25</i>
5.2.2	<i>Aligning Host I/O with RAID Striping.....</i>	<i>25</i>
	APPENDIX A: OBTAINING ADDITIONAL PERFORMANCE TOOLS	27
	APPENDIX B: SYSTEM TROUBLESHOOTING	28
	APPENDIX C: REFERENCES	29
	APPENDIX D: GLOSSARY OF TERMS	30

1 Audience and Scope

This document is intended to guide MD3000 and MD3000i customers through the advanced processes behind tuning their storage array to best suit their individual needs. Presented within are the best practices that must be observed in performance tuning a first generation firmware (06.XX.XX.XX) and second generation firmware (07.XX.XX.XX) storage array. For further detail on determining Firmware Generation of a MD3000 or MD3000i storage array, please view the Dell™ *Users Guide* at <http://support.dell.com/manuals>.

2 Performance Tuning Overview

The challenge of storage performance tuning is to understand and control the interacting components (listed below) while accurately measuring the application performance. Because the performance of the storage array accounts for only a portion of the overall application performance, tuning must be performed taking into account the input/output (I/O) characteristics of the application and all of the components involved in the data path, such as the SAS HBA, iSCSI initiator, network switch and host operating system settings.

With multiple aspects to consider, the task of performance tuning to maximize performance for even one application can seem formidable. Tuning the system to maximize performance of multiple applications that may share a single storage array can seem even more formidable. To reduce the complexity of tuning, Dell™ storage systems feature performance monitoring and flexible tuning controls that can be accessed through the Modular Disk Storage Manager (MDSM).

2.1 Components That Influence Storage Performance

This white paper provides an overall approach to tuning I/O performance and also provides specific guidelines for using the storage array tuning controls. These recommendations begin with an overall analysis of the elements that determine I/O performance:

- Storage array
- Application software
- Server platform (hardware, operating system, volume managers, device drivers)
- Network (for MD3000i only)

2.2 Basic Approach to Performance Tuning

The first principles of I/O performance tuning include the following question:

What should be the performance for my system?

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

The answers include:

- “It depends...” There are no absolute answers. Each environment is unique and the correct settings depend on the unique goals, configuration, and demands for the specific environment.
- “Actual mileage may vary.” Results vary widely because conditions vary widely.

The answers to this question suggest the following basic approach to performance tuning:

- 1 – Configure and test
- 2 – Measure
- 3 – Adjust as required

The performance monitoring features in all MD3000 and MD3000i storage systems and the controls for tuning make them ideally suited for this iterative process. The first step in tuning is to establish a baseline of existing performance. When generating a performance baseline, it is ideal to use a workload that is as similar to the intended final use of the storage solution. This can be as simple as the real application or a SQL Replay with a system performance monitor (perfmon or sysstat/iostat as well as the CLI and state-capture performance monitoring) or a synthetic benchmark package that closely imitates the expected I/O profile (Iometer, IOZone, Bonnie). By comparing the baseline data to the estimated needs and capability of the configuration, the user can effectively tune an MD3000 or MD3000i storage array. This white paper provides recommendations for this important first step as well as tuning tweaks to get the full capabilities from the MD3000 and MD3000i storage systems.

3 Application Software Considerations

In understanding the I/O characteristics of intended applications using the storage in a manner as close to the expected run-time is required to determine the optimal storage and system configuration, and is critical to tuning the overall solution. This includes, but is not limited to the following:

- Number of discrete I/O sources interacting with the solution
- Randomness of data access by the primary I/O source(s)
- Mean size of typical I/O; this is usually divided into the following three categories:
 - Large Block ($\geq 256\text{KiB}$) transfer size
 - Medium Block ($\geq 32\text{KiB}$ and $< 256\text{KiB}$) transfer sizes
 - Small Block ($< 32\text{KiB}$) transfer sizes
- Burstiness of I/O pattern, i.e. the mean duty-cycle of I/O to the storage array

- Profile of mean I/O direction; this is usually the ratio of Reads to Writes

4 Configuring the MD3000/MD3000i

There are two ways to configure the MD3000 and MD3000i storage systems. The most common and easiest method is using the MDSM. The MDSM allows automatic configuration settings that provide reasonable configurations with little knowledge of performance tuning required.

A manual configuration option is also available using the command line interface (CLI), providing more flexibility but requiring more knowledge of performance requirements.

For a link to the MDSM and CLI guides, see Appendix C: References.

4.1 *Determining the Best Raid Level*

The first step involved in tuning an MD3000 or MD3000i external storage array is to determine the RAID level most appropriate for the solutions given the application. In the proceeding document, please note that RAID 0 is excluded from most information in this paper due to the lack of data protection. This does not mean the use of RAID 0 is undesired, just that it should only be used for non-critical data. In general, RAID 0 generally provides better performance than RAID 1/10, 5 or 6. Additionally, RAID 6 is not always referenced specifically; most comments applicable to tuning RAID 5 are directly applicable to RAID 6 unless otherwise noted. In situations where the increased fault tolerance provided by RAID 6 are desired, please note that a performance penalty is observed when compared directly to RAID 5 due to the additional parity calculation and the extra physical disk required for implementation.

An important consideration when determining the appropriate RAID level is the physical disk cost required by a RAID level. Physical disk cost, is the number of physical drives worth of capacity that are sacrificed to provided the desired data integrity level. The physical disk cost of each RAID level is different, and may impact the decision for which RAID level is most appropriate for a chosen environment. RAID 0, having no level of redundancy has a physical disk cost of zero. RAID 1/10 has the highest disk-cost in disk groups containing more than 2 drives. Half of the physical drives in a RAID 1/10 are always consumed for the mirror. RAID 5 has a fixed 1-physical disk per disk-group cost, i.e. with a RAID 5 set of n disks, only n-1 worth of capacity is available. Similarly, RAID 6 has a fixed two physical disks per disk-group cost, or n-2. In RAID 5 and 6, these extra drives account for the space required to maintain the levels of parity information for each stripe.

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

Physical disk cost is not the only factor that influences the decision on which RAID level is most appropriate for a given application. The performance of a chosen RAID level is heavily interdependent on characteristics of the I/O pattern as transmitted to the storage array from the host(s). With I/O patterns involving write operations, when an I/O burst exceeds 1/3 of available cache memory in size, it should be considered a long I/O. Long writes show the performance of a chosen RAID level better than short writes. Short write operations can be handled entirely in cache, and the RAID level performance effect is minimized. As long as the write-burstiness is always lower than the cache to disk offload rate, a choice in RAID level can be a non-issue.

In general, the following outlines which RAID levels work best in specific circumstances:

- RAID 5 and RAID 6 works best for sequential, large I/Os (>256KiB)
- RAID 5 or RAID 1/10 for small I/Os (<32KiB)
- For I/O sizes in between, the RAID level is dictated by other application characteristics:
 - RAID 5 and RAID 1/10 have similar characteristics for most Read environments and sequential Writes.
 - RAID 5 and RAID 6 exhibit the worst performance mostly by random writes.
 - In random I/O applications consisting of more than 10% write operations, RAID 1/10 provides the best performance.

Table 1 provides a summary of these points for an ideal environment. An ideal environment consists of aligned stripes or segments reads and writes, as well as burst I/O where cache memory and RAID Controller Modules are not overly saturated by I/O operations.

Table 1: I/O Size and Optimal Raid Level

Block Size	Significantly Random		Significantly Sequential	
	Read	Write	Read	Write
Small (<32 KiB)	1/10, 5, 6	1/10	1/10, 5, 6	1/10, 5
Medium (Between 32 and 256 KiB)	1/10, 5, 6	1/10	1/10, 5, 6	5
Large (>256 KiB)	1/10, 5, 6	1/10	1/10, 5, 6	5

4.1.1 Selecting a RAID Level - High Write Mix Scenario

In random I/O applications with a >10% mix of write operations and a low degree of burstiness, RAID 1/10 provides the best overall performance for redundant disk groups.

RAID 1/10 performance can be >20% better than RAID 5 in these environments, but has the highest disk cost; for instance, more physical disks must be purchased. RAID 5 provides protection and minimizes disk cost for net capacity but is heavily impacted by the write performance overhead of parity updates.

RAID 6, while providing better protection than RAID 5 with minimal disk cost, is more heavily impacted by write performance overhead for the double parity calculations it requires.

In sequential I/O applications with relatively small write transfer sizes, RAID level does not make a large difference. With medium transfer sizes, RAID 1/10 can provide an advantage over RAID 5/6, again with a higher associated disk-cost. In very large sequential writes, RAID 5 can perform equal or better than RAID 1/10 especially when factoring in disk cost for equivalent capacity. Additionally, better performance is always achieved when the application or OS is capable of buffering or coalescing writes to fill a whole segment or stripe.

4.1.2 Selecting a RAID Level - Low Write Mix Scenario

In random I/O applications with a low mix (<10%) of write operations, RAID 5 offers approximately the same performance as RAID 1/10 in small transfer sizes but at a lower disk cost. RAID 0 provides somewhat better performance than 5 or 1/10 but offers no data protection. In larger transfer size environments, RAID 1/10 performances can be slightly better than RAID 5 but has a much higher disk cost.

4.2 Choosing the Number of Drives in a Disk Group

When optimizing performance, there are many factors to consider, such as drive type, capacity, and the number of drives.

The following general guidelines can be used in grouping drives in a disk group:

- Separate random and sequential workloads on different disks groups – segregating I/O traffic to minimize sharing of disk groups between virtual disks.
- Choose faster drives – in general, a single 15K RPM drive yields approximately 20% more performance than 10K for mixed sequential and random operations. Please refer to the manufacturer's data sheets to better determine the optimally suitable drive.
- Adding more drives to a disk group while holding the stripe size fixed can increase I/O rate for sequential I/O, up to the point of controller saturation – more drives mean more spindles to service I/O.

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

- To optimize for data transfer rate, multiply the number of physical data disks by the segment size to equal the I/O size. However, there are always exceptions. For small/medium I/Os care should be taken to avoid splitting up the I/O in such a way that an even smaller I/O is sent to the disk drives. Please note that data disks do not include parity or mirror disks used in a RAID set.

For IOPS or transaction-oriented applications, the number of drives becomes more significant because disk drive random I/O rates are relatively low. Select a number of drives that matches the per virtual disk group I/O rate needed to support the application. Make sure to account for the I/Os required to implement the data protection of the selected RAID level. Make the segment size at least as large as the typical application I/O size.

A segment size of 128K is a reasonable starting point for most applications. In most applications, the greater the number of drives in a disk group, the better the average performance. The drive count of an existing disk group can be increased using the CLI.

4.3 Virtual Disk Location and Capacity

The location of virtual disks within a disk group, the number and location of virtual disks allocated within a disk group, and the capacity of a virtual disk are important factors to consider when optimizing an array for performance.

When using rotating storage medium, the capacity of a virtual disk and its location within a disk-group greatly impact achieved performance. This is primarily due to differences in angular velocity in the outside zones. The effect of allocating the outermost edges of a rotational storage medium to increase performance is termed *short stroking* of a drive. While it is beyond the scope of this white paper to cover the technical details involved around short-stroking, typically the outer third of a rotational medium is the fastest, while the inner most zones are the slowest. Short stroking can easily be accomplished by creating a disk group consisting of a single virtual disk, which is allocated less than a third of the total capacity. The obvious downside to short-stroking a volume is the loss of additional usable capacity. Thus, this performance gain must be weighed directly against the capacity loss.

In addition to the performance gains of short-stroking, the effect of drive head seek should be taken into account when carving up a disk group into virtual disks. Virtual disks are aligned in series within a Disk Group with the first virtual disk starting in the fastest outer regions and progressing inwards. When taking this into account, a disk group should be designed with as few virtual disks as possible.

Dell™ does not recommend using more than four virtual disks or repositories per disk group for peak performance. Additionally, where performance is critical, isolate virtual disks to separate disk groups as possible. When multiple high-

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

traffic virtual disks share a disk group, even with purely sequential usage models, the disk group I/O behavior becomes increasingly random, lowering overall performance. Additionally, when a disk group must be shared, the most heavily trafficked virtual disk always should be located at the beginning of a disk group.


4.4 Virtual Disk Ownership

The Dell™ MDSM can be used to automatically build and view virtual disks. It uses optimal settings to stripe the disk group. Virtual disks are assigned to alternating RAID controllers when they are created. This default assignment provides a simple means for load balancing the workload of the RAID controllers. Ownership can later be modified to balance workload according to actual usage. If virtual disk ownership is not manually balanced, it is possible for one controller to have the majority of the work, while the other controller is idle.

Limit the number of virtual disks in a disk group. If multiple virtual disks are in a disk group, consider the following information:

- Consider the impact each virtual disk has on other virtual disks in the same disk group.
- Understand the patterns of usage for each virtual disk.
- Different virtual disks have higher usage at different times of day.

Figure 1: Virtual disk ownership balance.



4.5 Calculating Optimal Segment and Stripe Size

The choice of a segment size can have a major influence on performance in both IOPS and data transfer rate.

The term *segment size* refers to the amount of data written to one drive in a virtual disk group before writing data to the next drive in the virtual disk group. A set of contiguous segments spanning across member drives creates a *stripe*. For example, in a RAID 5, 4 + 1 virtual disk group with a segment size of 128KiB, the

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

first 128KiB of an I/O is written to the first drive, the next 128KiB to the next drive, and so on with a total stripe size of 512KiB. For a RAID 1, 2 + 2 virtual disk group, 128KiB would be written to each of the two drives (and same for the mirrored drives). If the I/O size is larger than this (the number of physical disks multiplied by a 128KiB segment), this pattern repeats until the entire I/O is complete.

For very large I/O requests, the optimal segment size for a RAID volume group is one that distributes a single host I/O across all data drives within a single stripe. The formula for maximal stripe size is as follows:

$$\text{LUN segment size} = \text{Maximal I/O Size} \div \text{number of data drives}$$

However, the LUN segment size should be rounded up to the nearest supported power of two value.

For RAID5 and 6, the number of data drives is equal to the number of drives in the volume group minus 1 and 2 respectively. For example:

$$\text{RAID5, 4+1 with a 64KiB segment size} \Rightarrow (5-1) \times 64\text{KiB} = 256\text{KiB stripe size}$$

Optimally, this RAID group is sufficient for handling I/O requests less than or equal to 256KiB.

For RAID1, the number of data drives is equal to the number of drives divided by 2. For example:

$$\text{RAID1/10, 2+2 with a 64KiB segment size} \Rightarrow (4-2) \times 64\text{KiB} = 128\text{KiB stripe size}$$

It is important to remember that depending on the application's I/O parameters, segment and strip size will vary.

For application profiles with small I/O requests, set the segment size large enough to minimize the number of segments (drives in the LUN) that are accessed to satisfy the I/O request, that is, to minimize segment boundary crossings. Unless specified otherwise by the application, starting with the default segment size of 128KiB is recommended.

It is imperative that the stripe size be correctly selected so that the host operating system is always making requests properly aligned with full stripes or full segments when possible.

4.6 Cache Settings

Read-ahead cache can be configured in the MDSM and through the CLI. The MDSM has only the defaults to work with, while the CLI can fully configure the read-ahead cache. Additionally, the global cache block size for read and write cache can be adjusted through the CLI.

Please refer to the *Dell™ PowerVault™ Modular Disk Storage Manager CLI Guide* on the Dell™ technical support site (<http://support.dell.com/manuals>) for a complete list of supported commands including the following cache specific commands.

4.6.1 Setting the Virtual Disk-Specific Write Cache and Write Cache Mirroring

Configured through the CLI – These commands are available at the virtual disk level.

Write Cache – Disabling write cache puts the controllers into a write-through mode, adding additional latency while data is flushed to the disk. Except for specific read-only environments, it is recommended that this setting stay enabled. Write cache is automatically disabled in the case of cache battery failure or a cache battery learn cycle.

Write Cache Mirroring – Write Cache Mirroring provides an additional level of redundancy and fault tolerance in the MD3000 and MD3000i. As a side effect, it reduces available physical memory and intra-controller bandwidth to perform this operation. In select, non data-critical cases, it can be beneficial to adjust this parameter. For normal use, Dell™ always recommends Enabling Cache Mirroring. Cache Mirroring is automatically disabled in the event of controller failure or when write caching is disabled.

WARNING: Data loss can occur if a RAID controller module fails while write-caching without cache mirroring is enabled on a virtual disk.

4.6.2 Setting the Virtual Disk-Specific Read Cache Pre-fetch

Configured through the CLI – This command is available at the virtual disk level.

Read Cache Pre-fetch – The Read Cache setting can be toggled on a virtual disk level. Disabling read pre-fetch is primarily useful in primarily small transfer size random read environments, where pre-fetch of random data would not provide sufficient value. However, the normal observed overhead of read pre-fetch is negligible. For most environments, Dell™ always recommends Enabling Read Pre-fetch.

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

4.6.3 Setting the Storage Array Cache Block Size

Configured through the CLI – This command is available at the storage array level and effects all virtual disks and disk-groups.

Cache Block Size – Cache Block Size refers to the way cache memory is segmented during allocation and affects all virtual disks in an array. On the MD3000 and MD3000i, settings of 4KiB and 16KiB are available with 4KiB being the default. A dramatic impact on performance can occur by choosing the correct cache block size setting specific to the system's I/O profile. If the typical I/O size is $\geq 16\text{KiB}$, which is typical with sequential I/O, set the storage array Cache Block Size to 16. For smaller ($\leq 8\text{KiB}$) I/O, especially in highly random or transactional use cases, the default 4KiB setting is preferred. As this setting affects all virtual disks on a storage array, changing it should be done with attention to the I/O needs of the application.

Table 2: Storage Array Standard Configuration Specifications

Option	MDSM GUI configuration templates			CLI Options
	File System	Database	Multimedia	
Drive Type	Selectable	Selectable	Selectable	Selectable
RAID Level	Selectable	Selectable	Selectable	0, 1/10, 5, 6
Segment Size ¹	128KiB	128KiB	256KiB	8KiB, 16KiB, 32KiB, 64KiB, 128KiB, 256KiB, 512KiB
Write cache with mirroring	Fixed to on	Fixed to on	Fixed to on	On or Off
Read-ahead cache	On	Off	On	On or Off
Cache Block Size ¹	Array Defaults to 4KiB			4KiB, 16KiB

¹ The MDSM CLI Guide and SMcli application may use the legacy terminology of KB to mean Kilobytes, or 2^{10} bytes. Throughout this whitepaper, the SI term kibibytes are used instead. When formulating an SMcli command the KB postfix is however still required. IEEE 1541-2002 and IEC 60027-2 standards detail units of measurement prefixes for binary multiples.

4.7 Tuning Using Array Performance Data

4.7.1 Collecting Performance Statistics

The stateCaptureData.txt and performanceStatistics.csv files, which are available through the MDSM, **Support** tab as part of a Technical Support Bundle, provide valuable statistical data in an easy-to-read format. The following section shows some sample data from the **stateCaptureData.txt** file and suggested configuration recommendations based on the performance considerations outlined in the previous section.

Other useful information is available through the array profile. Open the MDSM and select the Support tab – View Storage Array Profile.

Before collecting performance statistics, the I/O workload under test should be first executed. This will ensure the validity of performance statistics as part of the measure step of proper performance tuning.

Note: Figures shown below are from using the performance tool lometer.

4.7.2 RAID Level

The stateCaptureData.txt file provides statistics in the read and writes percentage columns to aid in selecting most appropriate RAID level. In Figure 2, the small reads and writes I/O percentages provide information regarding the distribution of the I/O types in the tested workload. This data is especially helpful when utilizing Table 1 referenced on page 7 to determining the applications current read/write mix. The RAID Level chosen can impact the I/O performance. Generally, RAID 1/10 has the best overall performance, with the highest physical disk cost. Use the I/O percent distribution and the average block size from the collected data to aid in this determination. These fields can be found in the highlighted regions of Figure 2 and Figure 3 for First and Second Generation Firmware respectively. It should be noted that the values in these figures are in block notation; block size for the specific virtual disk configuration is listed in the stateCaptureData.txt file, and is almost always 512 bytes. The average received I/O is not the I/O size the application uses but what the host sends, so while an application may attempt to send larger I/Os, the host's I/O stack may coalesce or break up I/Os as appropriate. Please see the appropriate separate OS or HBA documentation to determine these values.

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

Figure 2: First Generation Firmware - RAID Level. File: stateCaptureData.txt

```
Virtual Disk Unit 0 Configuration
Volume Type:          13+1 RAID 5
User Label:           MyRAID5_1
Block Size:           512 bytes
Large IO:             4096 blocks
Segment Size:         256 blocks
Stripe Size:         3328 blocks
...
IO Statistics:
      small      small      large      large      total      cache
      reads     writes     reads     writes     hits
requests 2028332119 147699066      0      0 2176031185 1289775370
blocks 3091968111 2518067526      0      0 1315068341 4019884678
avg blocks      4      17
IO pct.      93.21%    6.78%    0.00%    0.00%    0.00%    59.27%

      IOs      stripes      /IO      clusters      /IO
reads 2028332119 2034477363      1.00 2107869128      1.03
writes 147699066 148449472      1.00 157404718      1.06

write      Full      Partial      RMW No Parity      RMW2      FSWT
algorithms 1105611 12598366 32120072      0      0      0
```

Figure 3: Second Generation Firmware - RAID Level. File: stateCaptureData.txt

```
Volume 0 Attributes:
Volume Type:          RAIDVolume
User Label:           MyRAID10_One
...
BlockSize:           512 bytes
LargeIoSize:         4096 blocks
...
Perf. Stats:         Requests      Blocks      Avg. Blks      IO Percent
Reads                67456452 5943724625      88      71.20%
Writes                27283249 1144902648      41      28.80%
Large Reads           0      0      0      0.00%
Large Writes          0      0      0      0.00%
Total                 94739701 7088627273      74      100.00%
```

4.7.3 I/O Distribution

I/O can be characterized by its distribution and pattern. The two primary factors in determining the I/O distribution of an application are the randomness of I/O and the direction of I/O. The randomness of I/O indicates how sequential or random the data access is, as well as the patterning of this data access. The direction of the I/O can be simply related to the read and write percentages of I/O, that is, the direction I/O is taking from the storage device. I/O pattern refers

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

to how tightly the variance of sequential or random data access is contained within the volume. This can be purely random across an entire virtual disk, or random within some bounds, such as a large file stored within a virtual disk compared to large non-contiguous bursts of sequential data access randomly distributed within some bounds. Each of these is a different I/O pattern, and has a discrete case to be applied when tuning the storage.

The data from the stateCaptureData.txt file can be helpful in determining these characteristics. Sequential read percentage can be determined based on the percentage of total cache hits. If the cache hits and read percentage are high, then first assume the I/O pattern tends to be more sequential I/O. However, since cache hits are not broken out statistically by read and write, some variable experimentation may have to be performed with the representative data set if the pattern is unknown. For single threaded I/O host streams, this behavior can be confirmed by comparing the magnitude of reads to read pre-fetch statistics.

In cases where many sequential read operations are expected, enabling read pre-fetch in cache is recommended. If the cache hits percentage is low, the application tends to be more random and read-ahead should be disabled. Mid-range percentages possibly indicate bursts of sequential I/O but do not necessarily denote their affiliation to read or write I/O. Again, testing with read-ahead on/off would be required.

In the second generation firmware, the segment, stripe, and pre-fetch statistics have been reorganized as seen in Figure 4 from the lower half of Figure 2.

Figure 4: Second Generation Firmware - Performance Statistics Broken Out. File: stateCaptureData.txt

```
*** Performance stats ***
```

Cluster Reads	Cluster Writes	Stripe Reads
6252626	3015009	5334257
Stripe Writes	Cache Hits	Cache Hit Blks
2040493	4685032	737770040
RPA Requests	RPA Width	RPA Depth
982036	3932113	418860162
Full Writes	Partial Writes	RMW Writes
653386	29	328612
No Parity Writes	Fast Writes	Full Stripe WT
0	0	0

4.7.4 Stripe Size

For the best performance, stripe size should always be larger than the maximum I/O size performed by the host. As identified previously, stripes should be sized as even powers of two. The average block size can be identified from the collected data. Additionally, I/Os over 2MiB are considered large and broken out separately from smaller I/Os in the statistics. While all RAID level's benefit from careful tuning of stripe and segment size, RAID 5 and 6 with their parity calculations are the most dependant.

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

In the first generation of firmware (see Figure 5), this can be determined from the 'Avg. Blocks' row which represents the average I/O block size encountered. In Generation One, the 'Large IO' field denotes a 4096 block or 2MiB size with zero registered Large Reads or Writes during the sample period. Any single host received I/O greater than the Large I/O size is broken down into chunks less than or equal to the stated Large I/O value. It is extremely rare for a host to send such large I/Os.

With second generation firmware (see Figure 6), the 'Avg. Blks' column represents the average the I/O block size encountered. In Figure 6, the 'LargeloSize' field denotes a 2MiB size with zero registered Large Reads or Writes during the sample period.

Figure 5: First Generation Firmware - Stripe Size. File: stateCaptureData.txt

```
Virtual Disk Unit 0 Configuration
Volume Type:      13+1 RAID 5
User Label:       MyRAID5_1
Block Size:       512 bytes
Large IO:         4096 blocks
Segment Size:     256 blocks
Stripe Size:      3328 blocks
...
IO Statistics:
      small      small      large      large      total      cache
      reads     writes     reads     writes     hits
requests 2028332119 147699066 0          0 2176031185 1289775370
blocks   3091968111 2518067526 0          0 1315068341 4019884678
avg blocks 4          17         0          0 0          3
IO pct.  93.21%    6.78%     0.00%     0.00%    0.00%     59.27%
```

Figure 6: Second Generation Firmware - Volume Attributes Performance Stats; taken from RAID 1 volume. File: stateCaptureData.txt

```
Volume 0 Attributes:
Volume Type:      RAIDVolume
User Label:       MyRAID10_One
...
BlockSize:       512 bytes
LargeIoSize:     4096 blocks
...
Perf. Stats:
      Requests   Blocks   Avg. Blks   IO Percent
Reads  67456452  5943724625  88          71.20%
Writes 27283249  1144902648  41          28.80%
Large Reads 0          0          0          0.00%
Large Writes 0          0          0          0.00%
Total  94739701  7088627273  74          100.00%
```

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

In addition, the stateCaptureData.txt file provides a more granular method for determining the distribution of I/O within stripes and segments. In Figure 7 and Figure 8, item 1 is the number of full stripes read and written, and item 2 indicates the number of full clusters or segments read or written. The value of stripes per I/O request within reads and writes is also valuable for determining if stripe or segment settings are optimal for the tested data access pattern. The second generation firmware does not specifically break out the per-I/O ratio in data output as the first generation firmware did, however it can still be manually calculated simply by dividing the value in item 1 or 2 by the appropriate I/O request value from item 3 indicated in Figure 8.

In most cases, best performance is achieved with segment and stripe per-I/O ratios as close to 1.00. Traditionally, when tuning for maximal I/Os per second, if the segments per-I/O ratio is high, the current segment size may be too low for the application. Similarly, when tuning for highest possible data transfer rate, the stripe per-I/O ratio should ideally be 1.00 or even multiples. If this value is high, increasing the number of physical disks and/or the segment size can improve performance.

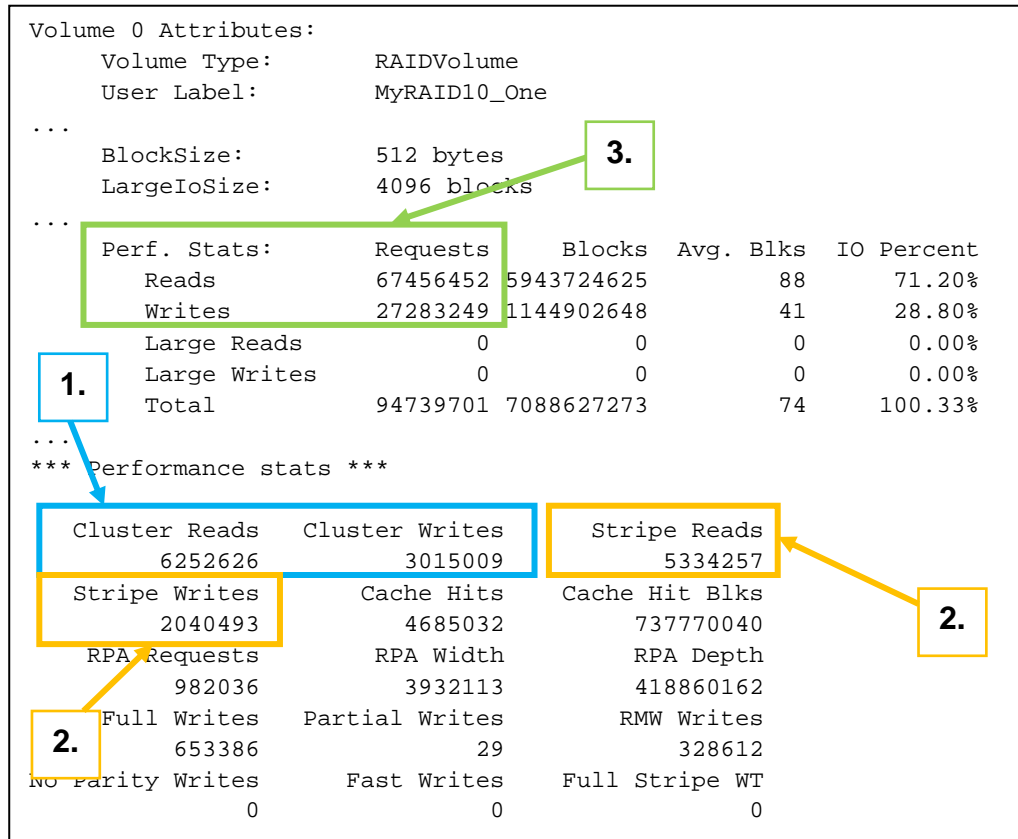
Figure 7: First Generation Firmware– Stripe distribution. File: stateCaptureData.txt

```
Virtual Disk Unit 0 Configuration
Volume Type:      13+1 RAID 5
User Label:       MyRAID5_1
Block Size:       512 bytes
Large IO:         4096 blocks
Segment Size:     256 blocks
Stripe Size:      3328 blocks
...
IO Statistics:
      small      small      large      large      total      cache
      reads     writes     reads     writes     hits
requests 2028332119 147699066      0      0 2176031185 1289775370
blocks 3091968111 2518067526      0      0 1315068341 4019884678
avg blocks      4      17      0      0      0      3
IO pct.      93.21%  6.78%  0.00%  0.00%  0.00%  59.27%

1.  IOs  stripes  /IO  clusters  /IO  2.
reads 2028332119 2034477363 1.00 2107869128 1.03
writes 147699066 148449472 1.00 157404718 1.06

write Full Partial RMW No Parity RMW2 FSWT
algorithms 1105611 12598366 32120072 0 0 0
```

Figure 8: Second Generation Firmware– Stripe distribution. File: stateCaptureData.txt



4.7.5 Write Algorithm Data

It is important to understand the effect determining the most suitable RAID level can be a daunting task. Understanding the effect of which write algorithm is in use is an important part of RAID Level balance. The possible choices in first generation firmware, as seen in Figure 9, are Full, Partial, RMW, RMW2, and the Full Stripe Write-Through. RMW2 was folded into RMW statistics in the second generation firmware (see Figure 10).

The Full algorithm takes an entire stripe of data and dumps it to the disk, depending on the RAID level of choice, P or P and Q will get calculated at this point. This is the most efficient type of write that can be performed, and the design of a disk group should be around maximizing full writes.

Partial writes are when less than a full stripe of data non-aligned with segment boundaries are modified and written. In RAID level 5 and 6 this is more complex, as parity data must be recalculated for the whole stripe. Partial writes are a worst-case algorithm and should be minimized. Larger quantities of partial writes than full writes can indicate use of an inappropriate segment size.

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

RMW, or Read-Modify-Write, is the second-best write algorithm available for RAID 5 and 6. A RMW occurs when a quantity of bits, smaller or equal to an individual segment are modified. This constitutes a two-read operation in RAID 5 and a three-read operation in RAID 6, with one of the segments being modified, and the parity drive(s) are read in. Then parity for the affected region is recalculated and then the data and parity in the stripe are re-written to disk. In small transactional processing, an extremely large number of RMWs should be expected. These RMW writes can cause a significant loss of performance; however this impact can be lessened by proper tuning of virtual disk stripe size.

RMW2 is used to differentiate Write-to-cache RMWs and Write-Through RMWs, with RMW2 being the latter, these statistics were consolidated in the second generation firmware. RMW2's also specifically happen when cache is disabled by force, failed mirroring controller (if policy is active) or failed cache battery. Additionally, the second generation firmware discretely tracks Full Stripe Write-Through conditions, and both generations track data on Number of Parity stripes re-calculated.

Figure 9: First Generation Firmware - Write Algorithms

```

Virtual Disk Unit 0 Configuration
  Volume Type:      13+1 RAID 5
  User Label:      MyRAID5_1
  Block Size:      512 bytes
  Large IO:        4096 blocks
  Segment Size:    256 blocks
  Stripe Size:     3328 blocks
  ...
  IO Statistics:
        small      small      large      large      cache
        reads     writes     reads     writes     total     hits
requests 2028332119 147699066      0          0 2176031185 1289775370
 blocks  3091968111 2518067526      0          0 1315068341 4019884678
avg blocks      4         17          0          0          0          3
 IO pct.   93.21%    6.78%    0.00%    0.00%    0.00%    59.27%

        IOs      stripes      /IO  clusters      /IO
 reads 2028332119 2034477363    1.00 2107869128    1.03
 writes 147699066 148449472    1.00 157404718    1.06

  write      Full      Partial      RMW  No Parity      RMW2      FSWT
algorithms 1105611 12598366 32120072      0          0          0
    
```

Figure 10: Second Generation Firmware - Write Algorithms

```

Volume 1 Attributes:
  Volume Type:      RAIDVolume
  User Label:      MyRAID5vd
  ...
  BlockSize:       512 bytes
  LargeIoSize:     4096 blocks
  ...
  Perf. Stats:
    Requests      Blocks  Avg. Blks  IO Percent
  Reads          577761063 1155647889 2      87.50%
  Writes         82542765 175687877 2      12.50%
  Large Reads    0          0          0      0.00%
  Large Writes   0          0          0      0.00%
  Total          660303828 1331335766 2      100.00%
  ...
  *** Performance stats ***

  Cluster Reads  Cluster Writes  Stripe Reads
  6252626        3015009        5334257
  Stripe Writes  Cache Hits     Cache Hit Blks
  2040493        4685032        737770040
  RPA Requests  RPA Width     RPA Depth
  982036         3932113        418860162
  Full Writes    Partial Writes  RMW Writes
  653386         29             328612
  No Parity Writes  Fast Writes    Full Stripe WT
  0              0              0
    
```

4.8 Using the CLI Performance Monitor

The CLI Performance Monitor is a command line scripting utility that provides access to all functions of PowerVault Storage Manager and some additional performance statistics. In Microsoft Windows, the scripting utility is named `smcli.exe` and is located by default in the `c:\program files\Dell\MD Storage Manage\client` directory. The complete reference guide *Command Line Interface Guide* for the MD3000/3000i is located on Dell's support website at <http://support.dell.com/manuals>.

The following are examples of CLI commands.

The command to run the **CLI Performance Monitor** is:

```

SMcli -n ArrayName -c "set session performanceMonitorInterval=5
performanceMonitorIterations=30;save storageArray performanceStats
file=\"performance.csv\";"
    
```

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

For a complete list of commands and instructions on using the CLI Performance Monitor, refer to the *Dell™ PowerVault™ Modular Disk Storage Manager CLI Guide* on <http://support.dell.com/manuals>.

4.9 Other Array Considerations

4.9.1 Global Media Scan Rate

Use the **Tools** tab to change/set the Media Scan settings in the MDSM. Global Media Scan uses CPU cycles and will affect performance if run at an inappropriate time, for example, during high-use user access or during backups. Set up and use Global Media Scan for low access logical drives, for example, file storage drives.

Please note that Dell™ does not recommend disabling media scan or lowering the media scan interval below 15 days. If Media-Scan is disabled, the risk of non-predicted failures can increase.

4.9.2 Setting the Virtual Disk-Specific Media Scan

Changing/setting the Media Scan settings in the MDSM is on the Tools tab. To choose to run Media Scan on specific virtual disk, highlight the appropriate virtual disk to be scanned, and select the **Scan selected virtual disks** check box.

4.10 Premium Feature Performance

4.10.1 Getting Optimal Performance from Snapshot

When distributing Snapshot Repositories, locate the repository virtual disks on separate drives from production virtual disks to isolate the repository writes and minimize the copy-on-write penalty. When possible, schedule Read I/Os to the Snapshot virtual disk to off-peak times when I/O activity on the source virtual disk is lower such as during evening hours.

4.10.2 Getting Optimal Performance from Virtual Disk Copy

The Virtual Disk Copy premium feature uses optimized large blocks to complete the copy as quickly as possible. Thus Virtual Disk Copy requires little tuning other than setting the copy priority to the highest level that still allows acceptable host I/O performance. Virtual Disk Copy performance is affected by other controller activity and by the RAID level and Virtual Disk parameters of the source virtual disk and the target virtual disk. A best practice for using Virtual Disk Copy is to disable all snapshot virtual disks associated with a source virtual disk before selecting the source virtual disk as a virtual disk copy target volume. Target and source virtual disks should ideally be resident on separate disk groups when possible; keeping them on the same disk group raises the potential for lower performing random I/Os for the copy operation.

5 Considering the Host Server(s)

5.1 Host Hardware Platform

5.1.1 Considering the Server Hardware Architecture

Available bandwidth depends on the server hardware. The number of buses adds to the aggregate bandwidth, but the number of HBAs sharing a single bus can throttle the bandwidth. Additionally, some server hardware has slower-speed PCI-E ports (4x) as well as high-speed ports (8x). The Dell SAS5e HBAs are PCI-E 8x devices and should be installed in 8x slots for maximum performance. Where additional PCI-E slots are available, two SAS HBAs should be used to redundantly attach the I/O host to each storage array controller module to maximize both performance and redundancy.

Note: Dell™ provides a bus layout on the lid of all servers. Consult this chart and use a different bus for each HBA installed in the Host.

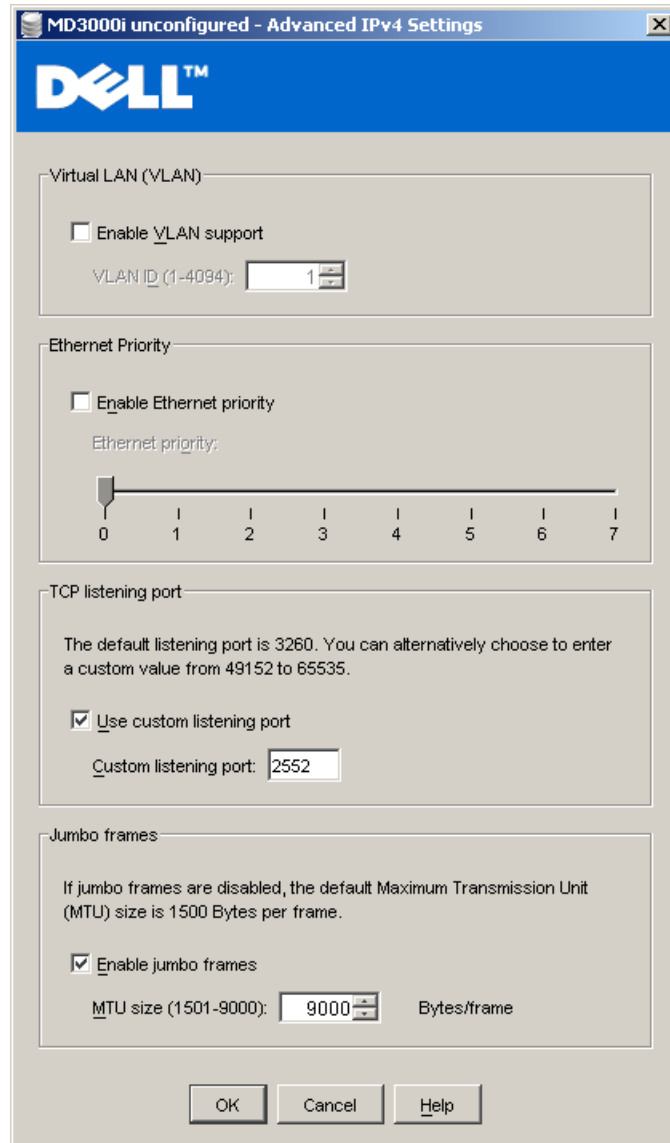
5.1.2 Sharing Bandwidth on the Dell™ MD3000i with Multiple NICs

Consider the following when sharing bandwidth on MD3000i with multiple NICs.

- Give each NIC its own IP address.
- Connect each NIC to a separate switch, or direct attach host NIC(s) to target iSCSI port(s).
- Use separate NICs for public network access than the storage array iSCSI session. Install additional NICs as needed.
- Ensure the use of separate redundant networks dedicated for iSCSI traffic, if this is not possible, setup a separate VLAN for iSCSI traffic.
- Use Jumbo Frames (Jumbo Frames increase the TCP frame size from 1500 bytes to 9000 bytes).
- The Microsoft iSCSI initiator will *not* work with teamed NICs.
- A single Host cannot mix HBAs and NICs to connect to the same or different arrays.

To edit the Jumbo Frames settings in the MDSM, select the **iSCSI tab, Configure iSCSI Host Ports, Advanced** (see Figure 11). Jumbo Frames can be also be set using the CLI. If using an MD3000i with Jumbo Frames, ensure the Ethernet switches and host NIC(s) also have Jumbo Frames enabled and are set to an equivalent value.

Figure 11: MD3000i Advanced IPv4 Settings: VLAN, QOS, Jumbo Frames Support



5.1.3 Sharing Bandwidth with Multiple SAS HBAs

Each SAS wide port includes four full duplex serial links within a single connector. The individual SAS 1.1 links run a maximum speed of 3Gb/s. A single path is used as the primary path to a connected device — the second, third, and fourth paths are used as overflow, when concurrent I/Os overload the primary channel. For example, if the first link is transmitting data at 3Gb/s, SAS uses 10-bit encoding versus 8-bit for byte transmission, which makes a single 3Gb/s capped at 300MiB/s. If another block of data then needs to be written to disk, for example, and link 1 is still busy, link 2 will manage the overflow of data that cannot be transmitted by link 1. If link 1 finishes its transmission of data, the next block of data will be transmitted on link 1 again, otherwise another link will be used. In this way, for heavy I/O workloads, it is possible that all links are being

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

used at certain times, providing a point-to-point raw wire speed of up to 12Gb/s. Please note that this raw speed does not take into account the transmission overhead or device expenditures on either side of the SAS wide link and is purely a cached I/O operation.

Additionally, care must be taken of which buses are being used within the host. Installing the HBAs that use the same bus will hamper data transfer rate. Ensure that all HBAs installed in the host are on a different bus (see note in 5.1)

5.2 Considering the System Software

5.2.1 Buffering the I/O

The type of I/O, buffered or un-buffered, provided by the operating system to the application is an important factor in analyzing storage performance issues. Un-buffered I/O (also known as *raw* or *direct* I/O) moves data directly between the application and the drive devices. Buffered I/O is a service provided by the operating system or by the file system. Buffering improves application performance by caching write data in a file system buffer, which the operating system or file system periodically flushes to non-volatile storage.

Buffered I/O is generally preferred for shorter and more frequent transfers. File system buffering might change the I/O patterns generated by the application. That is, writes might coalesce so that the pattern seen by the storage system is more sequential and more write-intensive than the application I/O itself. Direct I/O is preferred for larger, less frequent transfers and for applications that provide their own extensive buffering, such as Oracle. The I/O performance tool Iometer is another application that can operate un-buffered to test a more raw performance. Regardless of I/O type, I/O performance generally improves when the storage system is kept busy with a steady supply of I/O requests from the host application. Become familiar with the parameters that the operating system provides for controlling I/O, such as maximum transfer size.

5.2.2 Aligning Host I/O with RAID Striping

Most host operating systems require or benefit from varying degrees of partition I/O alignment and avoiding performance-degrading segment crossings. That is, I/Os should not span a segment boundary. Matching of I/O size (commonly, by a power-of-two) to disk group layout helps align I/Os across the entire drive. However, this is only true if the starting sector is properly aligned to a segment boundary. Segment crossing is often seen in the Microsoft Windows operating system, where partitions created by Microsoft Windows 2000 or Microsoft Windows 2003 start at the 64th sector. Starting at the 64th sector causes misalignment with the underlying RAID striping and allows the possibility for a single I/O operation to span multiple segments. Newer versions of the Microsoft Windows OS have a default alignment typically of 2048 blocks that still may require adjustment depending on the environment.

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

Microsoft provides the diskpar.exe utility as part of the Microsoft Windows 2000 Resource Kit, which was renamed to diskpart.exe in Microsoft Windows 2003 and later. Microsoft has a KB article 929491 covering this, and Dell™ always recommends checking for proper partition alignment to the stripe size of assigned virtual disks. Using these utilities, the starting sector in the master boot record can be set to a value that ensures sector alignment for all I/Os. Use a multiple of 64, such as 64 or 128. Applications such as Microsoft Exchange count on proper partition alignment to the disk stripe boundary.

For Microsoft's usage details about diskpart.exe, go to <http://technet.microsoft.com/en-us/library/aa995867.aspx>.

WARNING: *Changes made to the alignment of existing partitions **will** destroy data.*

Appendix A: Obtaining Additional Performance Tools

Table 3 shows a number of widely available tools, benchmarks, and utilities. Some of these tools are produced by non-profit organizations and are free.

Table 3: Performance Tools

Name	Description	Platform	Available From
IOBench	I/O data rate and fixed workload benchmark	Unix/Linux	http://www.acnc.com/benchmarks.html
Iometer	I/O subsystem measurement and characterization tool	Windows, Unix/Linux	http://www.iometer.org
IOZone	File system benchmark tool	Windows, Unix/Linux	http://www.iozone.org
Xdd	Tool for measuring and characterizing drive subsystem I/O	Windows, Unix/Linux	http://www.ioperformance.com
FIO	Unix/Linux benchmarking and I/O tool	Unix/Linux	http://freshmeat.net/projects/fio/
Bonnie++	Unix/Linux Filesystem I/O benchmark suite	Unix/Linux	http://www.coker.com.au/bonnie++/

Appendix B: System Troubleshooting

For information about troubleshooting MD3000 and MD3000i storage arrays, refer to the “Troubleshooting Problems” chapter of the *Dell™ PowerVault™ Modular Disk Storage Manager User’s Guide*. Go to:

MD3000:

<http://support.dell.com/support/edocs/systems/md3000/en/index.htm>

MD3000i:

<http://support.dell.com/support/edocs/systems/md3000i/en/index.htm>

Appendix C: References

Dell™ PowerVault™ Modular Disk Storage Manager CLI Guide,
<http://support.dell.com/support/edocs/systems/md3000/en/index.htm>

Dell™ PowerVault™ MD3000,
<http://support.dell.com/support/edocs/systems/md3000/en/index.htm>

PowerVault MD3000i SAN Array for Storage Consolidation,
http://www.dell.com/content/products/productdetails.aspx/pvaul_md3000i?c=us&1=en&s=bsd&cs=04

Using iSCSI: Dell™ PowerVault™ Modular Disk Storage Manager User's Guide,
<http://support.dell.com/support/edocs/systems/md3000/en/UG/HTML/iscsi.htm>

Dell iSCSI Cluster information,
http://www.dell.com/content/topics/global.aspx/sitelets/solutions/cluster_grid/clustering_ha?c=us&cs=555&l=en&s=biz&~page=3&~tab=4

Microsoft, using Diskpart,
<http://technet.microsoft.com/en-us/library/aa995867.aspx>

Microsoft, KB929491, Proper block alignment for NTFS using Diskpart,
<http://support.microsoft.com/kb/929491>

Microsoft, *Optimizing Storage for Exchange Server 2003*,
<http://www.microsoft.com/technet/prodtechnol/exchange/2003/library/optimizestorage.msp>

Dell High Availability Clustering: iSCSI
http://www.dell.com/content/topics/global.aspx/sitelets/solutions/cluster_grid/clustering_ha?&~page=3&~tab=4

IEC 60027-2 Ed. 2.0 (2000-11): SI Prefixes for Binary multiples
http://www.iec.ch/zone/si/si_bytes.htm

Appendix D: Glossary of Terms

Term	Definition
Burstiness	A data traffic property defined as the ratio of the peak I/O rate to the average I/O rate; in this case, the mean duty cycle exhibited by I/O transmitted or received from a storage array. Burstiness is adopted from its common usage in describing network workloads.
Controller Saturation	Controller saturation is where an individual RAID Controller Module is reached a maximum operational load and is unable to perform additional operations within its available bandwidth. Additional operations above this maximum point are not lost, they are instead queued up transitionally. This is considered an inflection point where performance hits a hard ceiling.
GiB	Gibibyte; see kibibyte.
HBA	Host Bus Adapter
HDD	Hard Disk Drive
Interposer	Integrated circuit with an interface routing between two opposing signaling sides. Used in the MD3000/MD3000i to interface SATA HDDs to the SAS backplane to perform translation.
iSCSI	Internet Small Computer Systems Interface. The iSCSI protocol is defined by the IETF in RFC 3720.
KiB	See Kibibyte.
Kibibyte	Kilo binary byte; IEC SI unit to separately denote 1024 or 2 ¹⁰ bytes (base-2) distinctly from the base-10 kilobyte (10 ³ or 1000 bytes). See IEC/ISO ISO 80000, IEC 60027-2, or IEEE 1541-2002.
Long I/O	Any I/O burst that exceeds 1/3 of the available cache memory size with an increased chance of not being wholly handled in cache.
MD3000	Dell™ PowerVault MD3000 Expandable Storage Array with SAS front-end.
MD3000i	Dell™ PowerVault MD3000i Expandable Storage Array with iSCSI front-end.
MDSM	Dell™ Modular Disk Storage Manager. Host management utility suite for configuring and maintaining a MD3000/MD3000i storage array.
MiB	Mebibyte, see kibibyte.
NIC	Network Interface Controller
NL-SAS	Near-line SAS; A hybrid technology of larger capacity SATA HDDs with a special SAS controller board

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

Term	Definition
RAID	directly on the drive in lieu of using a SATA interposer.
RAID 0	Redundant Array of Inexpensive Disks RAID Level 0; RAID 0 is a striped set with no redundant information. It is effectively a fully degraded RAID set with no disk redundancy overhead.
RAID 1/10	RAID Level 1/10: The RAID 1/10 implementation on the MD3000/MD3000i follows Berkley RAID 1 standard, expanding it to a redundant N+N mirrored set. Functionally, this implementation is set equivalent to a generic nested RAID 1+0 and operates on as few as 2 physical drives. This allows for any number of drive failures as long as one of each drive pair is available at the cost of half the included physical drives.
RAID 5	RAID Level 5; Involves a block-striping algorithm where n-1 disks per stripe in the raid set contain the data and a parity disk P contains parity or checksum used to validate data integrity provide consistency information. The parity is distributed across all drives in a disk group to provided further fault tolerance. RAID 5 provides protection from single drive failure.
RAID 6	RAID Level 6; Involves a block-striping algorithm where n-2 disks per stripe in the raid set contain the data and a parity blocks P and Q contains the parity or checksum used to validate data integrity provide consistency information. The parity is distributed across all drives in a disk group to provided further fault tolerance. RAID 6 provides protection from double drive failure.
RPA	Read Pre-fetch Algorithm: An acronym for the read ahead cache used on the MD3000\MD3000i.
RMW	Read, Modify, Write; The second best algorithm available in RAID 5 and RAID 6 write operations. A RMW occurs when a quantity of bits, smaller or equal to an individual segment are modified.
RMW2	A Firmware Generation One adaptation of RMW specifically for Write-Through conditions where write cache is not enabled or requested to a virtual disk.
SAS	Serial Attached SCSI. Protocol maintained by t10.org
SATA	Serial Attached Technology Attachment, or Serial ATA. This is the next stage of the legacy Parallel ATA. Primarily in this white paper SATA refers to SATA HDD technology.
Saturation	See Controller Saturation
SCSI	Small Computer System Interface, Protocol maintained

Dell™ PowerVault MD3000 and MD3000i Array Tuning Best Practices

Term	Definition
	by t10.org
Segment	A segment is the data written to one drive in a virtual disk group stripe before writing data to the next drive in the virtual disk group stripe.
Short I/O	Any I/O that consumes less than 1/3 of available cache memory that can be handled within cache, effectively a cached operation.
SQL	Structured Query Language; flexible markup language for computer databases maintained by ANSI and ISO.
Stripe	A set of contiguous segments spanning across member drives.