

## **Solution Guide – Choosing Storage for Dell Database Solutions with Dell™ PowerVault™ and Dell/EMC Storage**

### **Abstract**

This white paper provides a guide for choosing an appropriate storage architecture for Dell Database Solutions deployed on Dell PowerEdge servers with Dell PowerVault and Dell/EMC storage. Using the knowledge gained through joint development, testing and support with Microsoft and Oracle, this white paper documents “best practices” that can help you select a storage solution to meet your capacity, performance and availability needs.

October 2007

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.

© 2007 Dell Inc. All rights reserved.

Reproduction in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

Trademarks used in this text: *Dell*, the *DELL* logo, *PowerEdge* and *PowerVault* are trademarks of Dell Inc.; *Intel* and *Xeon* are registered trademarks of Intel Corporation; *EMC*, *Navisphere*, and *PowerPath* are registered trademarks of EMC Corporation; *Microsoft*, *Windows*, and *Windows Server* are registered trademarks of Microsoft Corporation.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

October 2007 Rev. A01

# Table of Contents

<b>INTRODUCTION .....</b>	<b>4</b>
DELL DATABASE SOLUTIONS FOR ORACLE 10G AND MICROSOFT SQL SERVER 2005 .....	4
OVERVIEW OF THIS WHITE PAPER .....	4
<b>OVERVIEW OF DELL STORAGE SYSTEMS .....</b>	<b>5</b>
DIRECT-ATTACHED SERIAL-ATTACHED SCSI (SAS).....	5
<i>PowerVault MD3000 Direct-Attached Storage Array</i> .....	5
<i>MD3000 Capacity</i> .....	5
<i>MD3000 Availability</i> .....	6
INTERNET SCSI (iSCSI) STORAGE NETWORKING .....	7
<i>PowerVault MD3000i iSCSI Storage Array</i> .....	7
<i>MD3000i Capacity and Availability</i> .....	8
FIBRE CHANNEL STORAGE NETWORKING .....	8
<i>Dell/EMC CX3 Series FC Storage Arrays</i> .....	8
<i>Dell/EMC CX3 Series Capacity</i> .....	9
<i>Dell/EMC CX3 Series Availability</i> .....	9
<b>COMMON CRITERIA FOR CHOOSING A STORAGE ARRAY.....</b>	<b>11</b>
ARCHITECTURAL COMPLEXITY .....	11
COST .....	11
SCALABILITY.....	12
<b>OVERVIEW OF DATABASE PERFORMANCE WITH DELL STORAGE ARRAYS.....</b>	<b>13</b>
<b>SELECTING A STORAGE ARRAY FOR A DELL DATABASE SOLUTION .....</b>	<b>15</b>
<b>CONCLUSIONS.....</b>	<b>17</b>
<b>TABLES AND FIGURES INDEX .....</b>	<b>18</b>
<b>REFERENCES .....</b>	<b>18</b>

## Introduction

Dell™ PowerEdge™ servers, Dell PowerVault™ storage systems and Dell/EMC storage systems are ideal choices to deploy highly reliable and sustainable Oracle® Database 10g and Microsoft® SQL Server™ 2005 solutions. This white paper is intended to help IT professionals design and configure Oracle 10g and SQL Server 2005 database solutions using Dell servers and storage that apply “best practices” derived from laboratory and real-world experiences. This white paper provides criteria and guidance to assist in the selection of the appropriate storage system for your Dell Database Solution.

### *Dell Database Solutions for Oracle 10g and Microsoft SQL Server 2005*

Dell Database Solutions for Oracle 10g and Microsoft SQL Server 2005 are designed to simplify operations, improve utilization and cost-effectively scale as your needs grow over time. In addition to providing price/performance leading server and storage hardware, Dell Solutions for Oracle 10g and SQL Server 2005 include:

**Dell Tested and Validated Configurations** – in-depth testing of Oracle 10g and SQL Server 2005 configurations with Dell servers and storage; documentation and tools that help simplify deployment  
**Integrated Solution Management** – standards-based management of Dell Database Solutions that lower operational costs through integrated hardware and software deployment, monitoring and update  
**Database Software Licensing** - multiple licensing options that can simplify customer purchase  
**Dell Enterprise Support and Professional Services for Dell Database Solutions** – offerings for the planning, deployment and maintenance of Dell Solutions for Oracle 10g and SQL Server 2005

Dell PowerEdge servers and Dell Storage help to minimize operating costs with price/performance leadership - Dell currently holds price/performance leadership for TPC-E and seven of the top ten TPC-C price/performance leadership positions with SQL<sup>1</sup>.

For more information about Dell Solutions for SQL Server 2005, see [www.dell.com/sql](http://www.dell.com/sql).

For more information about Dell Solutions for Oracle 10g Database, see [www.dell.com/oracle](http://www.dell.com/oracle).

### *Overview of this White Paper*

When deploying a database solution, IT professionals are faced with a range of choices for storing their data, each choice providing varying capacities, performance and availability features. This white paper is intended to guide the reader with a set of criteria and recommendations that can be used to identify the appropriate Dell Storage system for your Dell Database Solution.

---

<sup>1</sup> Source: TPC-E by Price/Performance Version 1 and TPC-C by Price/Performance Version 5 Results as of October 2007. See [www.tpc.org](http://www.tpc.org) for current results

## Overview of Dell Storage Systems

Dell offers several different storage options that can be deployed into database solutions. Modern storage systems generally employ one or more of the following technologies: Serial-Attached SCSI (SAS), Internet SCSI (iSCSI), and Fibre Channel (FC). Each of these technologies has a shared heritage, as they have each evolved to overcome the limitations and constraints of the traditional parallel SCSI infrastructure. All three technologies were designed with advanced storage array features, such as multi-path I/O and multi-initiator environments (e.g. clusters), in mind.

This section provides an overview of each of these technologies and the application of the technologies in current Dell storage arrays that help make them well suited for use with database solutions. For database solutions that feature a single, non-clustered, database server, any of the storage arrays discussed in this paper may be connected directly to the host server. To better support the deployment of clustered database servers using Microsoft Cluster Service (MSCS) with SQL Server 2005 or Oracle Real Application Clusters with Oracle Database 10g, these storage arrays also provide storage-based RAID engines which provide write cache mirroring and battery backup systems for the cache memory. These features allow storage array write caching to be enabled, even in multi-initiator environments, a function generally not possible with server-based RAID controllers.

### *Direct-Attached Serial-Attached SCSI (SAS)*


Direct-attached storage refers to a topology in which a storage system and server(s) are connected to one another without any switches or storage network components. This topology can be employed with most existing storage device interconnect technologies, including Serial-Attached SCSI (SAS), Fibre Channel (FC), and Internet SCSI (iSCSI). Direct-attached storage topologies not only offer a simple and cost-effective way to increase the storage capacity of a particular host, but can also provide storage-based RAID and the ability to connect multiple hosts – features that were once only available with more-complex networked storage topologies, including network attached storage (NAS) and storage area network (SAN).

### **PowerVault MD3000 Direct-Attached Storage Array**

The Dell PowerVault MD3000 is a high-performance storage array designed to be connected via a direct-attached topology to servers that are equipped with SAS host bus adapters, such as the Dell SAS/5E. Figure 1 provides an example of connecting a single Dell PowerEdge server to the MD3000. The MD3000 features dual-ported SAS hard-disk drives and is connected to the host servers using SAS wide-link (x4) cables. In a direct-attached topology, the number of servers which can connect to the storage is determined by the number of host ports available on each RAID controller, and whether hosts are configured to use redundant connections.

### **MD3000 Capacity**

Each MD3000 unit supports up to 15 SAS drives and can be expanded by daisy-chaining up to two MD1000 storage arrays for a total of 45 SAS drives. The drives provide capacities of up to 300 gigabytes and either 10,000 RPM or 15,000 RPM speeds. Each of the four SAS ports available on the MD3000 operate at a peak of 3.0 gigabits per second (Gbps) for an aggregate peak throughput of 12 Gbps.




**Figure 1 - Single Server Directly Attached to a PowerVault MD3000 Storage Array**

### **MD3000 Availability**

Although the MD3000 can be operated with a single RAID controller, Dell best practices recommend that redundant controllers be used to increase the resiliency of the solution and to maximize the availability of the data stored by a database solution. Each RAID controller on the MD3000 features two host ports. Therefore, it is possible to connect up to four hosts, each with a single data path – however, this introduces single points of failure in the storage subsystem, and is not optimal for database solutions. The Dell recommended configuration for database solutions features an MD3000 with redundant RAID controllers and provides for the connection of two hosts, each with redundant data paths (e.g. one connection from each server to each RAID controller). This configuration shown in Figure 2 greatly helps reduce the possibility that a single point of failure might cause any connected server to lose access to the storage system. I/O paths into RAID controllers are managed by a multi-path driver, which re-routes I/O request from the hosts to the other controller when a failure of one of the I/O paths occurs. To further enhance availability, each RAID controller includes a backup battery that can hold cached data in the controller’s cache memory for up to 72 hours without external power. Redundancy is provided for cached writes as they are mirrored between controllers.

Additional availability and integrity features of the MD3000 include hot- pluggable, redundant power supplies, cooling modules and disk drives; active disk scrubbing and non-disruptive firmware upgrades. In addition, the MD3000 offers two optional Premium Features: Snapshot Virtual Disks and Virtual Disk Copy. Snapshot Virtual Disk creates point-in-time snapshots of a data volume. Using a copy-on-first-write technique, data on source virtual disks can be updated without changing the contents of the point-in-time snapshot. Snapshots are often used as part of a backup system, reducing the time during which a database must be in a quiescent state. Virtual Disk Copy produces a full block-by-block copy of a source virtual disk, which can be utilized in the same manner as “clones” or “business continuity volumes”, such as for disk-based backups, recovery to a state preserved via a snapshot, or migration of data to a larger volume. For more information about these features, refer to the Dell PowerVault Modular Disk Storage Manager User’s Guide, which is available from <http://support.dell.com/>.



**Figure 2 - Cluster Directly Attached to a PowerVault MD3000 Storage Array**

### ***Internet SCSI (iSCSI) Storage Networking***

Internet SCSI (iSCSI) is an industry standard that encapsulates SCSI block I/O commands and data for transmission via a TCP/IP network. The iSCSI protocol can be used to establish and manage connections between storage devices and hosts using standard Ethernet networks. iSCSI storage devices may be directly-attached, or can be deployed as networked solutions, which are often called IP SANs. Each host in an iSCSI solution is configured with iSCSI initiators that connect to iSCSI targets on the storage device. These iSCSI ports are interconnected by means of standard Gigabit Ethernet hardware. In order to communicate with an iSCSI storage device, each attached host must employ an iSCSI initiator. Software iSCSI initiators are available for Microsoft® Windows Server® and Linux operating systems, and allow common GbE adapters to be used in lieu of dedicated storage host bus adapters (HBAs). GbE adapters that feature a TCP/IP Offload Engine (TOE), can also be used, but have been shown to have a minimal impact to the overall performance of the storage subsystem - typically less than a 5% improvement in iSCSI throughput.

### **PowerVault MD3000i iSCSI Storage Array**

The Dell PowerVault MD3000i is designed as a flexible IP SAN storage array that employs the iSCSI protocol for communication with host servers. Like the MD3000, the MD3000i features dual-ported SAS hard-disk drives, and is offered with a single controller which provides two 1 Gb/s Ethernet (GbE) ports. Dell best practices for database solutions recommends that each MD3000i storage array use dual RAID controllers for redundancy, with each controller featuring two GbE ports for a total of four 1Gb Ethernet ports for server host connections. In this configuration, the MD3000i can be directly-attached to two hosts, and each host has an available data path to each RAID controller to provide redundancy in the event of a component failure. The MD3000i and the accessing servers can also be interconnected via GbE switches. In this configuration, shown in Figure 3, the MD3000i may provide storage for up to 16 separate or clustered hosts.

## MD3000i Capacity and Availability

Since the PowerVault MD3000i and MD3000 share a common RAID controller architecture, the MD3000i provides the same maximum storage capacity of 45 drives and implements a common availability feature set, including support for the Snapshot Virtual Disks and Virtual Disk Copy Premium Features.




Figure 3 - Dell PowerVault MD3000i IP SAN Architectural Overview

## Fibre Channel Storage Networking

Fibre Channel (FC) is a multi-gigabit speed network technology that is primarily used for storage networking, providing a high-performance interconnect for servers and storage systems that can operate of long distances with low latency and low transmission error rates. FC networks support various industry standard protocols including SCSI and can be implemented with direct connect and switched fabric topologies. Supporting up to hundreds of end nodes with operating speeds of up to 4 Gbps, and inter-node distances of up to 2 kilometers<sup>2</sup>, FC networks are well suited for implementing large-scale Storage Area Networks (SAN) providing high performance and high availability.

## Dell/EMC CX3 Series FC Storage Arrays

The Dell/EMC CX3-series storage arrays are designed as scalable and flexible Fibre Channel (FC) storage arrays, featuring dual-ported 4 Gb/s FC hard-disk drives and 4 Gb/s FC host connections. The series includes the CX3-10c, CX3-20, CX3-40, and CX3-80 Storage arrays. The differences between these arrays are highlighted in Table 1, and include the following: number and speed of the CPU(s) in the storage processors (SPs), the number of host connections (front-end ports), the number of hard-disk loop connections (back-end ports), the size of the SP read/write cache, and the number of supported hard-disk drives.

<sup>2</sup> The maximum distance between components is dependent on the speed and on type of optical ports used. Refer to the documentation for the storage array, FC HBA, and FC switch in your solution for more details.



## Dell/EMC CX3 Series Capacity

Depending on which model is selected, the CX3-series supports anywhere from 60 to 480 hard disks attached to the same array, and when deployed in a SAN up to four arrays can be utilized by the same host server. A single CX3-80 array can support over 350 TB of raw disk storage. The specifics for each model in the CX3-series are detailed in Table 1. The front-end ports listed in the table are used either for direct-attached hosts or to connect the CX3-series array to an FC SAN. In either case, redundant connections are recommended to provide multi-path I/O and to ensure continued data availability for database servers in the event of many types of failures. The CX3 series architecture is shown in Figure 4.

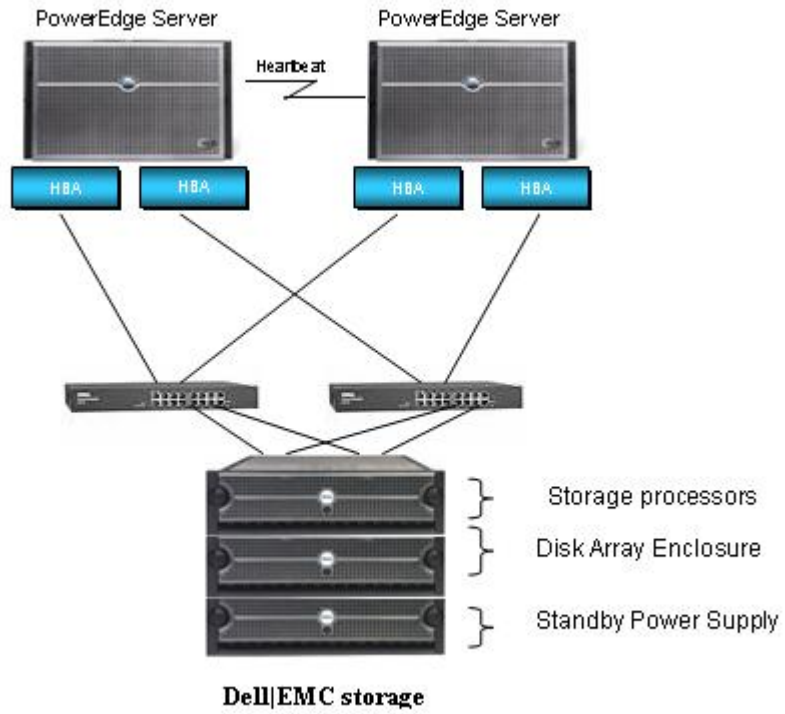
Features	CX3-10C	CX3-20 CX3-20c	CX3-20f	CX3-40 CX3-40c	CX3-40f	CX3-80
No. of CPUs per SP	1	1	1	2	2	2
CPU Frequency	2.8 GHz	2.8 GHz	2.8 GHz	2.8 GHz	2.8 GHz	3.6 GHz
Array Cache (2 SPs)	2 GB	4 GB	4 GB	8 GB	8 GB	16 GB
Maximum Drives per Array	60	120	120	240	240	480
Maximum Additional DAEs per Array	3	7	7	15	15	31
Maximum No. of redundantly-attached hosts	64	128	128	128	128	256
4 Gbps FC front-end ports per array (2 SPs)	4	4	12	4	8	8
4 Gbps back-end loops (2 SPs provide redundancy)	1	1	1	2	4	4

**Table 1 – Dell/EMC CX3 Series Comparison**

## Dell/EMC CX3 Series Availability

The Dell/EMC CX3 storage systems are made up of the following modular components: storage processor enclosure (SPE), disk array enclosure (DAE) and dual standby power supplies (SPS). The SPE contains a pair of storage processors (SP) to provide redundancy. Each DAE houses up to 15 FC or SATA disks. For primary storage for database solutions, Dell recommends the use of FC disks. A Dell/EMC CX3-series storage array requires at least one DAE, which contains the operating software for the storage array. Additional DAEs may be added to increase the capacity of the storage system. The SPS enables proper shutdown of the storage system during power failure by powering the SPE and first DAE long enough to allow the SPs to safely move the data currently in the write cache to a vault area on the first DAE. When configured with two SPs, the CX3-series also provides a mirrored write cache.

All Dell/EMC CX3 series arrays also support optional features, including SnapView snapshots and clones, MirrorView/Asynchronous, MirrorView/Synchronous, and SANCopy. SnapView snapshots and clones offer functions similar to the Virtual Snapshot Disk and Virtual Disk Copy Premium Features on the PowerVault MD3000 and MD3000i. MirrorView allows mirroring of a LUN from one CX3-series array to another array, and SANCopy allows copying or migrating LUNs between separate arrays.



**Figure 4 - Dell/EMC CX3 Series Architectural Overview**

## Common Criteria for Choosing a Storage Array

There are several factors to consider when selecting a storage system for a database solution. Since database solutions fill a wide range of needs and are deployed into vastly different customer environments, the key factors that influence this decision vary from case to case. It is also important that the storage system is considered in the context of the entire database solution, including not only the storage hardware, but also servers, interconnect components, and software. The Dell storage systems discussed above will be compared in terms of the following criteria: architectural complexity, storage system cost, and scalability. Additionally, certain performance criteria for these storage systems will be discussed. Refer to the “Overview of Database Performance with Dell Storage Arrays” section for this analysis.

### *Architectural Complexity*

The MD3000 and MD3000i share a common architecture, except for the ports used to connect host servers. Therefore, it is appropriate to compare these systems collectively to the Dell/EMC storage arrays. The RAID controllers in the PowerVault MD-series storage systems serve as the “brains” of the storage array. These controllers perform the RAID calculations, control the I/O movement, communicate with the management client, store the firmware, and protect data until it can be written safely to the hard disk drives. MD-series storage arrays offer redundant RAID controllers with internal cache backup batteries, power supplies, cooling fans, and hard disk drives in a single three rack unit (3U) enclosure. Dell PowerVault MD1000 storage systems, which provide more disk capacity, each require an additional 3U. Conversely, the Dell/EMC CX3-series storage systems feature separate modules for storage processors, disk drives and standby power supplies. A CX3-series storage system requires at least 5U for the first 15 hard-disk drives. Additional DAEs each require 3U. The integrated design of the Dell PowerVault MD-series is easier to understand and simpler to deploy.

In addition, different storage protocols and topologies impose different physical distance limitations among these Dell storage systems. The MD3000, by nature of employing a direct-attached topology, requires that the servers are physically close to the storage system, and is determined by the maximum allowable length of a SAS x4 cable, which is approximately four meters. Dell/EMC Fibre Channel SAN and MD3000 IP SAN allow greater distance between the storage and server, allowing more flexible deployment within a datacenter. Distance-extension solutions for Fibre Channel and iSCSI may be used to enable replication of data over considerable distances between sites as part of a business continuity or disaster recovery plan.

### *Cost*

The components required by a Fibre channel storage array such as the Dell/EMC CX3-series, including the server HBA, optical cabling, switches and other infrastructure components tend to be more expensive than the equivalent components required for other storage solutions. For example, iSCSI can significantly lower the cost of implementing networked storage by employing lower-cost GbE technology, and allowing re-use of existing infrastructure investments. Direct-attached storage has even lower costs as it does not require a switching infrastructure. SAS storage systems, such as the PowerVault MD3000 offer many of the benefits of networked storage without the additional expense of the storage network infrastructure. However, if future growth and the ability to scale by adding additional storage – beyond what is supported with a single storage system – are required, deploying a SAN architecture may be beneficial. Both Dell/EMC CX3-series and MD3000i iSCSI storage arrays offer this flexibility.

Along with the cost in hardware, deploying and managing FC SAN requires specialized expertise, which is obtained by the investment in training personnel or hiring a service company. IP SANs may provide a cost savings alternative in terms of personnel and training. Because Ethernet networking is a well established technology and is well understood by most IT staff, little to no additional network training may be needed to enable IT staff to set up an IP storage network. Direct-attached storage requires no networking expertise.

## *Scalability*

The Dell storage systems discussed in this paper vary in terms of total storage capacity and number of hosts which can be connected, both key considerations when evaluating current and future requirements for a database solution. The total storage capacity of the storage system is a function of the size and number of hard-disk drives offered by the solution, which is – in turn – a function of the number of arrays that can be connected to a host server and the number of expansion enclosures that can be attached to each array. A database solution can only employ a single MD3000 array, which supports up to two MD1000 expansion enclosures for a total of 45 total hard-disk drives. Although currently available drive capacities as large as 400Gbytes can support large databases, database performance is more typically determined by the number of drives over which the data is stored, not their capacity. When a database deployment is deemed to require more than 45 total disks, a solution with multiple MD3000i or CX3-series arrays can be deployed. Both of these solutions offer the ability to attach up to four storage arrays to the same host server(s) in a SAN environment. With four MD3000i arrays, each with two MD1000 expansion enclosures, a total of 180 disks can be made available for a database solution. With four maximally-configured CX3-80 arrays, a total of 1920 disks can be deployed; refer to Table 1 to determine the maximum number of disks for each CX3-series array.

When building Oracle RAC clusters or deploying SQL Server 2005 into failover clusters, the number of hosts than can be connected to a given array while maintaining redundant connections between each server and the storage array is an important consideration. The MD3000 allows the connection of up to two host servers with redundant paths. The MD3000i allows the connection of up to two host servers with redundant paths when configured in a direct-attached topology, or up to 16 host servers when connected via Ethernet switches in an IP SAN. The CX3-series provides for between two and six directly-attached servers, depending on the number of front-end, or host, ports on each SP, and up to 256 initiator ports (or 128 redundantly-attached servers) when deployed in an FC SAN.

Clearly, networked storage solutions offer greater scalability, both in terms of total data capacity and of host connectivity. If a database solution has high initial requirements, then the limitations of a direct-attached solution may prove unsuitable. Similarly, if the solution has modest initial requirements, but is expected to grow substantially in the future, an iSCSI or FC solution can be directly-attached initially, deferring the cost of interconnect, and later be migrated into a SAN when additional storage capacity or host connections become a necessity.

## Overview of Database Performance with Dell Storage Arrays

The performance of the storage subsystem is a key consideration when planning a database deployment and performance should be considered in terms of the small random I/O and large sequential I/O workloads typical of databases. Using database benchmarks and I/O generation tools, the performance of different storage architectures can be measured for the classic database workloads - online transaction processing (OLTP) and query-based business intelligence systems most often referred to as data warehousing (DW). OLTP applications typically generate small (8KB) random reads and writes. The throughput in I/O operations per second (IOPS) and the I/O response time are the key criteria for comparing the performance of storage subsystems for OLTP databases. On the other hand, query-based DW systems tend to generate sequential read and write streams composed of multiple outstanding large (1 MB) I/Os. The bandwidth of a storage subsystem, in Megabytes per second (MB/s), is the key criterion for DW configurations.

To measure the storage systems discussed in this paper, the Dell Database Solutions Engineering team employed an array of database benchmarks and I/O generation tools, including the following: IOMeter<sup>3</sup>, Quest Software® Benchmark Factory® TPC-C, the Oracle I/O Numbers (ORION) tool, and the Microsoft BenchCraft kit for TPC-E. The I/O-generation tools IOMeter and ORION do not require the installation of a database; they operate by exercising the storage subsystem consistent with OLTP and DW database workloads. In addition, Oracle Database 10g and Microsoft SQL Server 2005 were deployed for specific database benchmark runs. With each of these tools, care was taken to isolate the storage subsystem, and comparisons were made with each tool using identical server, physical disk, RAID level, and LUN / virtual disk configurations. Different test iterations exercised each storage system with configurations ranging from four to forty-five disk drives. The results generated by each tool were aggregated and normalized to determine the suitability of the storage architectures for various database workloads. The analysis of these results is summarized in the following paragraphs and in Table 2.

For OLTP workloads, characterized by a preponderance of small random writes, both I/O-generation tools and database benchmark tools consistently demonstrated that SAS, iSCSI, and Fibre Channel arrays provide roughly equivalent throughput (IOPS) results. When maintaining an I/O response time that is consistent with industry “best practices” for Oracle Database 10g or Microsoft SQL Server 2005, the PowerVault MD3000i iSCSI array significantly outperformed the PowerVault MD3000 SAS and Dell/EMC CX3 fibre channel arrays and generated higher throughput when a small number of disks (up to 12) were used. As the number of disks was increased, this advantage was eroded, and the I/O latency on the iSCSI array increased more rapidly than was observed with either the SAS or Fibre Channel arrays. For configurations with more than 20 disks, the MD3000 and CX3-series were shown to produce OLTP database benchmark results, in terms of transactions per second (TPS) that were 10-20% higher than those produced by the MD3000i. These results are useful for the many small databases that are deployed in production environments. Also, when applying best practices for an Oracle 10g database, a flash recovery area (FRA) that is twice as large as the database storage area is recommended. Hence, the number of disks used for the primary database storage remains consistent with the range in which the MD3000i performs well.

For DW workloads, characterized by a preponderance of large sequential reads, results were consistent regardless of the number of disk drives tested. The I/O-generation tools consistently demonstrated that the SAS and Fibre Channel arrays delivered higher MB/s throughput than the iSCSI arrays. Although the theoretical per host port bandwidth of an MD3000 SAS wide (x4) port exceeds that of a 4 Gbps Fibre Channel port, the CX3 arrays maintained a slight performance edge over the MD3000. This difference is likely to be attributed to the difference in read-ahead caching on the arrays. In these tests, the performance of the MD3000i was limited as a result of its GbE host ports, which offer the lowest theoretical and observed bandwidth of the arrays mentioned in this paper. Although the theoretical maximum throughput of the MD3000 and CX3 storage arrays are up to twelve times higher, the observed performance of the MD3000i was between 25-35% less than its peers.

---

<sup>3</sup> Open source measurement and characterization project; for details see <http://www.iometer.org/>  
Choosing Storage for Dell Database Solutions

	<b>Transaction Processing</b>	<b>Data Warehousing</b>
<b>Dell PowerVault MD3000</b>  <b>Direct-Attached SAS Array</b>	<p>The MD3000 array is capable of handling the throughput (IOPS) required for OLTP workloads. This array is a good choice for implementations that fall within its scalability limits. However, for database clusters of more than two nodes, or for database solution implementations that require more than 45 disk drives, a networked storage array should be chosen.</p>	<p>The MD3000 array is capable of handling the bandwidth (MB/s) required for DW workloads. This array is a good choice for smaller DW implementations; however, the primary limitation of the MD3000 array is its scalability – both in terms of the number of supported hosts, and of the maximum possible database size.</p>
<b>Dell PowerVault MD3000i</b>  <b>Direct- or IP SAN-Attached iSCSI Array</b>	<p>The MD3000i is capable of handling the throughput (IOPS) required for OLTP workloads. In fact, for small OLTP databases (20 or fewer disk drives), this array can produce up to 150% higher throughput than the other arrays reviewed in this paper. For larger databases - up to 45 disk drives - the MD3000i provides 10-20% lower throughput than the other arrays. When additional scalability is required, up to four MD3000i arrays can be connected to the same database host or database cluster.</p>	<p>The MD3000i array is less suitable for DW workloads. The primary limitation of this array is its GbE iSCSI interfaces, which result in 25-35% lower bandwidth (MB/s) than is observed with the other arrays reviewed in this paper.</p>
<b>Dell/EMC CX3-series</b>  <b>Direct- or SAN-Attached Fibre Channel Arrays</b>	<p>The CX3-series arrays are capable of handling the throughput (IOPS) required for OLTP workloads. Their increased data cache size, and support for single volumes larger than 2 TB combine to make these arrays a strong choice for the largest OLTP implementations.</p>	<p>The CX3-series arrays are capable of handling the bandwidth (MB/s) required for DW workloads. Their increased data cache size, support for a greater number of connected hosts, and higher capacity limits combine to make these arrays a good choice for larger DW implementations.</p>

**Table 2 - Suitability of Dell Storage Arrays for Common Database Workloads**

## Selecting a Storage Array for a Dell Database Solution

The Dell storage arrays discussed in this whitepaper feature architectures and designs with differing theoretical performance maximums. However, real-world performance is influenced by more than the host interfaces and interconnection technologies used by the storage arrays, and it is not appropriate to recommend a particular storage array based on a single criterion. As this paper has illustrated, there are a number of considerations that should be carefully weighed when selecting a storage array for a particular database solution. Common selection criteria include the following: price, performance, suitability for a given workload, scalability of disk capacity and of host connectivity, and storage array feature set particularly high availability features. Figure 5 provides a flowchart illustrating a possible decision making process that weighs these criteria to help determine the Dell storage array that is best suited for the needs of an intended database deployment.

The process illustrated in Figure 5 shows how certain solution design criteria may influence the decision to choose direct-attached versus networked storage, or to choose Fibre Channel versus iSCSI for networked storage. The research conducted for this whitepaper have helped shape the following general “rules of thumb” for choosing a storage array:

- Price/Performance remains a key consideration. Fibre Channel arrays – which provide the greatest available capacity and the highest absolute performance – are the most costly
- If more than two hosts are required for a database solution, then networked storage topologies should be selected
- When future capacity needs are projected to greatly exceed the current needs, an iSCSI or Fibre Channel array can be initially configured in a DAS topology, and later migrated into a SAN
- If required storage capacity (current or planned) exceeds what can be provided by a single DAS storage array, then networked storage should be selected
- When a small number of disk drives are used, the drives will likely become a performance constraint before the storage array host interface becomes a bottleneck. The available bandwidth of a particular host interface takes on increased significance as additional disk drives are added to the configuration
- For OLTP workloads – which primarily feature small random write operations – SAS, iSCSI, and Fibre Channel storage arrays are all capable of providing suitable throughput (IOPS), making cost and scalability the more likely decision points
- For DW workloads – which primarily feature large read operations – the higher bandwidth host interfaces of the SAS and Fibre Channel arrays provide more bandwidth (MB/s) than is attainable with current GbE iSCSI arrays.

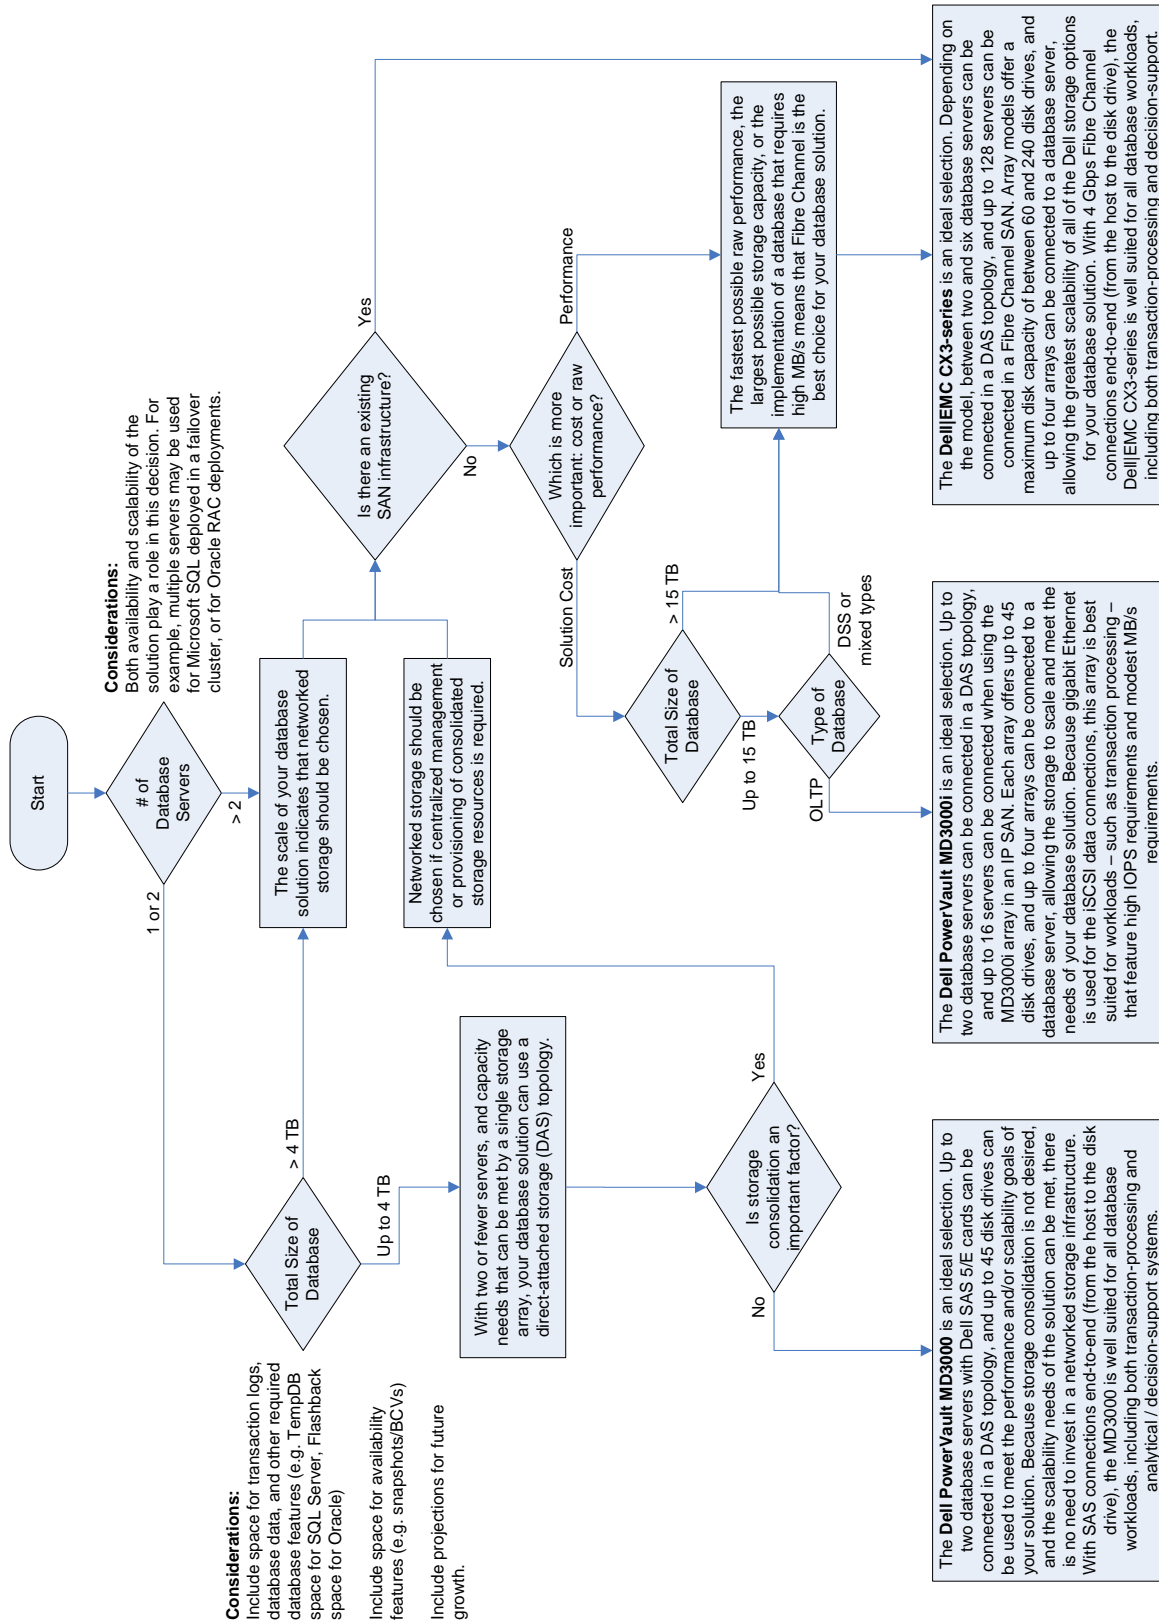


Figure 5 - Flowchart for Selecting a Storage Array for a Dell Database Solution



## Conclusions

Dell Database Solutions are designed to simplify operations, improve utilization and cost-effectively scale as your needs grow over time. This white paper provides guidelines that can help you select the appropriate Dell Storage array for your database needs.

After evaluating multiple factors, including scalability, cost, and performance of various database workloads, there are some notable differences between the various storage architectures discussed in this paper. Each array has specific advantages with respect to these criteria that must be evaluated before selecting the optimal array for a particular database solution. By examining the results presented in this paper, and using a logical process to evaluate each criterion, a good fit can be found for a variety of sizes and types of database deployments.

The Dell PowerVault MD3000, with its SAS x4 wide host interfaces, provides a range of availability features and offers throughput and bandwidth suitable for many database workloads. This storage array supports up to two redundantly-connected hosts and a total of 45 disk drives - scalability that can meet most entry and many mid-range database needs.

The Dell PowerVault MD3000i, with its GbE iSCSI host interfaces, provides a range of availability and scalability features and offers throughput suitable for OLTP database workloads. Up to 16 redundantly-connected hosts may be attached via an IP SAN, and up to four arrays – each supporting 45 disk drives – may be connected to a single host. The bandwidth available from the MD3000i host interfaces make it more suitable for OLTP and less suitable for mid-range to large business intelligence solutions.

The Dell/EMC CX3-series, with their 4 Gbps Fibre Channel host interfaces, provide a range of availability features and the best scalability features of the arrays discussed in this paper. The CX3 series offer throughput and bandwidth suitable for demanding OLTP and business intelligence workloads.

The best practices described here are intended to help achieve optimal performance of Oracle Database 10g and SQL Server 2005 on PowerEdge servers and Dell storage. To learn more about Dell Database Solutions, please visit [www.dell.com/oracle](http://www.dell.com/oracle) and [www.dell.com/sql](http://www.dell.com/sql) or contact your Dell representative for up to date information about Dell servers, storage and services for Dell Database solutions.

## Tables and Figures Index

Table 1 – Dell/EMC CX3 Series Comparison.....	9
Table 2 - Suitability of Dell Storage Arrays for Common Database Workloads.....	14
Figure 1 - Single Server Directly Attached to a PowerVault MD3000 Storage Array.....	6
Figure 2 - Cluster Directly Attached to a PowerVault MD3000 Storage Array.....	7
Figure 3 - Dell PowerVault MD3000i IP SAN Architectural Overview.....	8
Figure 4 - Dell/EMC CX3 Series Architectural Overview.....	10
Figure 5 - Flowchart for Selecting a Storage Array for a Dell Database Solution.....	16

## References

1. “Pro Oracle Database 10g RAC on Linux”, Julian Dyke and Steve Shaw, Apress, 2006.
2. “Benchmark Factory for Databases”, Quest Software.  
[http://www.quest.com/Quest\\_Site\\_Assets/PDF/Benchmark\\_Factory\\_5\\_TPCB.pdf](http://www.quest.com/Quest_Site_Assets/PDF/Benchmark_Factory_5_TPCB.pdf)
3. “Interface Decisions – FC, iSCSI, and SAS”, Dell Inc.  
[http://www.dell.com/downloads/global/products/pvaul/en/iscsi\\_fc\\_sas\\_interface.pdf](http://www.dell.com/downloads/global/products/pvaul/en/iscsi_fc_sas_interface.pdf)
4. “PowerVault MD3000i brochure”, Dell Inc.  
[http://www.dell.com/downloads/global/products/pvaul/en/pvaul\\_md3000i\\_brochure.pdf](http://www.dell.com/downloads/global/products/pvaul/en/pvaul_md3000i_brochure.pdf)
5. “PowerVault MD3000 brochure”, Dell Inc.  
[http://www.dell.com/downloads/global/products/pvaul/en/pvaul\\_md3000\\_brochure.pdf](http://www.dell.com/downloads/global/products/pvaul/en/pvaul_md3000_brochure.pdf)
6. “Dell/EMC CX storage Products Brochure”, Dell Inc.  
[http://www.dell.com/downloads/global/products/pvaul/en/dlemc\\_broch.pdf](http://www.dell.com/downloads/global/products/pvaul/en/dlemc_broch.pdf)