

Dell™ PowerVault™ Solutions for Microsoft® SQL Server™ 2005 Always On Technologies

Abstract

This technical note provides information on the Dell™ PowerVault™ storage solutions, based on the Microsoft® SQL Server™ Always On Storage Solution Review program. Under the program, Dell has documented compliance for the Dell PowerVault MD3000 and MD3000i storage-array products.

September, 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. *Oracle* is a registered trademark of Oracle Corporation and/or its affiliates.

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.

September 2007 Rev. A01

Table of Contents

ABSTRACT	1
OVERVIEW.....	4
DISCLAIMER	4
FEATURES.....	4
DELL COMPLIANCE FOR SQL ALWAYS ON	4
CORE 1.00: WINDOWS LOGO CERTIFICATION	5
CORE 1.01: WINDOWS API SUPPORT	5
CORE 1.02: STABLE MEDIA.....	5
CORE 1.03: FORCED UNIT ACCESS AND FORCED WRITE-THOUGH	5
<i>Battery Backup</i>	5
<i>Mirrored Write Cache</i>	5
CORE 1.04: ASYNCHRONOUS CAPABILITIES	5
CORE 1.05: WRITE ORDERING.....	6
CORE 1.06: TORN PAGE PROTECTION.....	6
CORE 1.07: NTFS SUPPORT.....	6
RECOMMENDED 1.08: PARTNER ENHANCED ESCALATION	6
APPENDIX A: MICROSOFT ALWAYS ON SPECIFICATION VERSION 1.02.....	7
<i>Requirements</i>	7
<i>List of core requirements</i>	7
APPENDIX B: REFERENCES	10
<i>General reference</i>	10
<i>Dell PowerVault customer documentation</i>	10

Overview

This tech note provides information on the Dell™ PowerVault™ storage solutions, based on the Microsoft® SQL Server™ Always On Storage Solution Review program.

Information regarding the Microsoft SQL Server 2005 Always On program is available at <http://www.microsoft.com/sql/alwayson>.

Disclaimer

This document is produced independently of Microsoft Corporation. The information contained in this document represents the current view of Dell on the issues discussed as of the publication date. Because of changing market conditions, it should not be interpreted as a commitment on the part of Dell. Also, Dell cannot guarantee the accuracy of any information presented after the publication date.

Features

The Microsoft SQL Server Always On program provides formal compliance documentation for solutions provided by vendors. Dell provides a number of highly available solutions, which are covered by this framework. This document provides coverage of those storage-based solutions, which are compliant under the definition of Microsoft SQL Server Always On.

Dell Compliance for SQL Always On

The subsequent sections document Dell compliance for the requirements as provided by SQL Server Always On documentation. The specifications that apply to this Always On submission are listed in Appendix A: Microsoft Always On Specification Version 1.02.

Under the program, Dell has provided compliance with the Dell PowerVault MD3000 and MD3000i storage-array products and any associated PowerVault MD1000 storage expansion units.

Core 1.00: Windows Logo certification

All Dell PowerVault storage platforms listed in the “Dell Compliance” section are logo certified. Windows logo certification requirements and platform listings are available in the Windows Catalog under the storage array category.

The following URL provides a link to the Windows Catalog. Search “Dell” under the storage category.

<http://www.microsoft.com/windows/catalog/server/>

Dell additionally provides solutions for clustering with PowerVault arrays. Microsoft provides a listing within the Window Catalog for logo certification under the **Hardware classification > Cluster Solutions**.

Core 1.01: Windows API support

Dell PowerVault storage platforms making up the storage arrays under this compliance statement fully support the core Windows API.

Write operations to supported PowerVault storage platforms guarantee delivery to stable media as defined in subsequent sections of this document. Cache write operations are protected by battery backup systems and other cache protection mechanisms, such as cache write mirroring.

Core 1.02: Stable media

All Dell PowerVault storage arrays covered under this compliance statement fully adhere to SQL Server Write Ahead Logging (WAL) protocols and meet ACID (Atomicity, Consistency, Isolation, and Durability) requirements as defined in the SQL Server 2000 I/O Basics documentation. PowerVault storage array replication features ensure that log predecessor writes are honored. These solutions utilize consistency technology.

Core 1.03: Forced Unit Access and Forced Write-through

All Dell PowerVault storage arrays under this compliance statement adhere to Forced Unit Access and Forced Write-through requirements.

PowerVault storage arrays use onboard caching mechanisms to optimize I/O operations for connected servers and associated applications. Write operations specifically benefit from the speed of write operations to cache. Cache I/O operations are typically orders of magnitude faster than write operations to the physical disk media. PowerVault storage arrays utilize a protection mechanism to help ensure the durability and persistence of updated (write) data stored within the cache.

Battery Backup

Dell PowerVault storage arrays include battery backup components. In the event of a failure in the primary power supply, data in the cache can be maintained up to 72 hours via battery power.

When primary power is restored, any pending updates stored in the cache memory are submitted to the relevant logical units. In no case are partial I/O operations propagated to the logical unit.

Mirrored Write Cache

Dell PowerVault arrays implement write cache mirroring protection. Under this scheme, updated cache areas are implemented in a RAID 1 convention. As a result of this implementation, updates are fully redundant and are protected against a single point of failure, such as a memory board fault.

Core 1.04: Asynchronous Capabilities

The Dell PowerVault storage arrays under this compliance statement adhere to this requirement. PowerVault storage arrays will not transition asynchronous I/O operations from a host into synchronous operations.

Core 1.05: Write Ordering

The Dell PowerVault storage arrays under this compliance statement adhere to and can enforce write ordering.

Core 1.06: Torn Page Protection

The Dell PowerVault storage arrays under this compliance statement support this core requirement.

Core 1.07: NTFS Support

The Dell PowerVault storage arrays covered under this compliance statement provide full support for all NTFS capabilities.

Recommended 1.08: Partner Enhanced Escalation

Dell is a Microsoft Gold Certified Support partner and both companies have signed a Cooperative Support Agreement (CSA). The agreement formally commits Dell and Microsoft to work together for mutual customers with valid configurations and support contracts. Additionally, the agreement defines call escalation and transfer processes that make support resources available to each company 24x7x365 on a worldwide basis.

Appendix A: Microsoft Always On Specification Version 1.02

Requirements

This section contains the core requirements for the Always On Storage Solution Review program as provided by Microsoft when this compliance document is submitted. Storage system capabilities are divided into two categories: Required and Recommended.

Type	Definition
Required	A capability or property the subsystem <i>must</i> provide to pass the requirements of the Always On Storage Solution Review Program.
Recommended	A capability or property the subsystem <i>should</i> provide for optimal compatibility with SQL Server.

An Always On solution white paper must document the product, feature or features and specific product configurations that support each of these core requirements before you submit the paper for review to the Always On Storage Solution Review program.

List of core requirements

Core 1.00 Windows Logo certification	Required
Microsoft Windows logo certification helps ensure the safety of Microsoft SQL Server data by testing various aspects of the products. To be compliant with the Always On Storage Solution Review program, solutions must pass the Certified for Windows logo testing.	

Core 1.01 Core Windows API support	Required
Microsoft SQL Server uses several APIs to accomplish secure data storage. A storage solution must ensure that a system supports specific API properties throughout the various layers and implementations of the I/O solution.	

Core 1.02 Stable media	Required
Microsoft SQL Server relies on the Write-Ahead Logging (WAL) protocol to maintain the Atomicity, Consistency, Isolation, and Durability (ACID) properties of the database and to guarantee data integrity. WAL relies on stable media capabilities. A solution must comply with this stable media guarantee. For detailed information, consult “Power Outage Testing – Pull The Plug” section in Chapter 2 of the <i>Microsoft SQL Server I/O Basics</i> .	

Core 1.03 Forced Unit Access (FUA) and Write-Through	Required
To support Write-Ahead Logging (WAL), Microsoft SQL Server uses FILE_FLAG_WRITETHROUGH and FlushFileBuffers to open files. Both of these options must be supported by storage solutions. All components in a solution must honor the write-to-stable media intent. This includes, but is not limited to, caching components. It is not enough to only honor WAL for Microsoft SQL Server log files. Data files and backup streams also depend on WAL behavior. Some storage products include battery-backed caching mechanisms as part of the write-through guarantee. If these caching mechanisms are present in the solution, the Always On solution white paper should document the practical limits of the write-through guarantee for a production environment.	

For more information, consult the links listed in the References section at the end of this paper, and the following Microsoft Knowledge Base article, KB917043 - Key factors to consider when you evaluate third-party file cache systems with Microsoft SQL Server.

Core 1.04 Asynchronous capabilities	Required
<p>Microsoft SQL Server performs most of its I/O using asynchronous capabilities. If a request specifies asynchronous operation, no API call should cause a synchronous condition. Synchronous I/O can cause unexpected scheduler and concurrency issues. Therefore, an Always On solution must provide asynchronous I/O capabilities.</p> <p>For more information, consult the white paper, <i>How To Diagnosis and Correct Errors 17883, 17884, 17887, and 17888</i> at: http://www.microsoft.com/technet/prodtechnol/sql/2005/diagandcorrecterrs.msp</p>	

Core 1.05 Write ordering	Required
<p>A tenet of the WAL protocol is write ordering or order preservation. An Always On solution must provide write ordering capabilities.</p> <p>The write-ordering requirement applies to both local and remote I/O destinations. If a database is split among physical paths, all paths must honor the ordering across all files related to database. To satisfy this ordering requirement, products sometimes use a user-defined consistency group that encapsulates all the database files. An Always On solution white paper must include information about the configuration requirements needed for the solution to meet the write ordering requirement. For example, a solution requiring a consistency group might specify this configuration requirement as: “All files associated with a database must be configured in a single consistency group.”</p> <p>Example Configuration</p> <p>A solution has the following configuration:</p> <ul style="list-style-type: none"> • Data File A LUN —Subsystem #1 • Log File B LUN—Subsystem #2 <p>If these subsystems use separate physical paths with different caching, Microsoft SQL Server is unable to support this configuration because the caching mechanisms may not present a coherent cache. The subsystems would require a third element to maintain cache coherency across the disparate caches.</p> <p>The same caching problem described in the example configuration can also occur across network boundaries. If a database backup is written to a UNC path but FlushFileBuffers only guarantees that the local system file cache is flushed, Microsoft SQL Server is exposed to data loss.</p> <p>For more information about write ordering requirements, consult the “Write Ordering,” “FlushFileBuffers,” “Backup Hardening,” and “Remote Mirroring” sections of the white paper, <i>Microsoft SQL Server 2000 I/O Basics</i>.</p>	

Core 1.06 Torn I/O protection	Required
<p>An Always On solution must provide sector alignment and sizing in a way that prevents torn I/O. This includes splitting I/Os across various I/O entities in the I/O path.</p> <p>Additionally, a solution must accurately report sector size to Windows I/O APIs. Accurately reporting sector size helps prevent sector size mismatches and avoid torn writes. For example, a drive that does 4 KB writes reports 512 bytes while the drive performs a read/write of the 4 KB sectors. This inaccuracy in reporting sector size can create a condition where data is lost and exposed as torn writes. An Always On solution must document configurations in such a way that use sector sizes from the sector size list that is supported by Microsoft SQL Server : 512, 1024, 2048, and 4096 bytes.</p> <p>To indicate when a torn-write situation occurs, it is recommended that the solution log appropriate warning events.</p> <p>The Always On solution white paper must include information about the configuration requirements needed for the solution to meet the torn I/O requirements.</p> <p>For more information, consult the “Torn I/O,” “Log Parity,” and “Sector Size” sections located in Chapter 2 of the white paper, <i>Microsoft SQL Server I/O Basics</i>.</p>	

Core 1.07 NTFS support	Required
<p>You must support NTFS capabilities. This includes, but is not limited to, the following:</p> <ul style="list-style-type: none"> • Sparse Files • File Streams • Encryption • Compression • All Security Properties <p>You must support sparse files on NTFS based file systems. Microsoft SQL Server 2005 uses sparse files in support of DBCC CHECK* commands and snapshot databases.</p> <p>Common copy and compression utilities may not honor sparse file metadata, but instead copy all bytes, ignoring the sparse allocations, and requiring full storage space. Vendors may choose to provide utilities to copy or move sparse files without destroying the sparse file intent.</p> <p>Support must be provided for streams on NTFS based files. Microsoft SQL Server 2005 uses sparse file streams files in support of DBCC CHECK* commands.</p>	

Core 1.08 Partner Enhanced Escalation Process membership	Recommended
<p>The Partner Enhanced Escalation Process (PEEP) is a joint escalation process established between Microsoft and partners. The non-binding memorandum (MOU) establishes direct, cross company joint escalation assistance and issue management.</p>	

Appendix B: References

General reference

[Microsoft SQL Server 2000 I/O Basics](#) -

<http://www.microsoft.com/technet/prodtechnol/sql/2000/maintain/sqlIObasics.msp> - applies to SQL Server 7.0, SQL Server 2000, and SQL Server 2005

Dell PowerVault customer documentation

Dell provides documentation for Dell PowerVault storage arrays through the Dell support website (<http://support.dell.com>)

For more information on Dell PowerVault storage arrays and integration with Microsoft SQL Server consult www.dell.com/sql .