

**Technical Note**

P/N 300-003-798

REV A01

June 8, 2006

---

This technical note contains information on these topics:

- ◆ Overview ..... 3
- ◆ Disclaimer ..... 3
- ◆ Features ..... 3
- ◆ EMC Compliance ..... 3
- ◆ Core 1.00 – Windows Logo Certification ..... 4
- ◆ Core 1.01 – Windows API Support ..... 4
- ◆ Core 1.02 – Stable Media ..... 4
- ◆ Core 1.03 – Forced Unit Access and Forced Write-through ..... 5
- ◆ Core 1.04 – Asynchronous Capabilities ..... 6
- ◆ Core 1.05 – Write Ordering ..... 6
- ◆ Core 1.06 – Torn Page Protection ..... 11
- ◆ Core 1.07 – NTFS Support ..... 13
- ◆ Recommended 1.08 – Partner Enhanced Escalation ..... 13
- ◆ Appendix A – Microsoft Always On Specification Ver 1.02 ..... 14
- ◆ Appendix B – References ..... 19

Copyright © 2006 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

## Overview

This document provides information on the EMC® Symmetrix® 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 has been produced independently of Microsoft Corporation. Microsoft Corporation expressly disclaims responsibility for, and makes no warranty—express or implied—with respect to the accuracy of the contents of this document.

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

## Features

The Microsoft SQL Server Always On program provides formal compliance documentation for solutions provided by vendors. EMC 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 SQL Server Always On.

## EMC compliance

The subsequent sections document EMC 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 ver 1.02."

Under the program, EMC has provided compliance with the Symmetrix DMX™ storage array series of products.

EMC Symmetrix DMX series: Engenuity™ microcode 5670 and later in combination with EMC Solutions Enabler 6.2 or later.

## Core 1.00 – Windows Logo certification

All EMC storage platforms listed in the “EMC compliance” section are Logo Certified. Windows logo certification and platform listings are available on the Windows Catalog under the storage array category.

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

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

EMC additionally provides a solution in the Geographically Dispersed clustering listing. Microsoft provides a listing within the Window Catalog for logo certification under the Hardware classification > Cluster Solutions > Geographically Dispersed Cluster Solutions. Within the Geographically Dispersed Cluster Solutions, search for “SRDF/CE.”

The EMC SRDF<sup>®</sup>/CE for Microsoft Cluster Service product utilizes SRDF functionality as documented within this compliance document. As such, SRDF/CE for Microsoft Cluster Service also complies with the Microsoft SQL Server 2005 Always On requirements for a supported storage platform.

## Core 1.01 – Windows API support

EMC storage platforms comprising the storage arrays under this compliance statement fully support the core Windows API.

Write operations to supported EMC 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 and cache destaging to “cache vault.”

## Core 1.02 – Stable media

All EMC 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. EMC storage arrays and replication products ensure that log predecessor writes are honored. These solutions utilize EMC consistency technology.

## Core 1.03 – Forced Unit Access and Forced Write-through

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

EMC storage arrays are integrated cache disk arrays (ICDA). These systems provide 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.

All EMC storage arrays utilize a protection mechanism to ensure the durability and persistence of updated (write) data stored within the cache. Specifically, for storage arrays included within the Always On program, a number of mechanisms are provided.

### Battery backup

EMC storage arrays include battery backup devices. These battery backup solutions are tested and certified to support the required operations in the event of a failure in the primary power supply.

For DMX-3 arrays under primary power failure, cache memory is written to persistent durable media in a designated cache vault located on specific disks within the array. When primary power is restored, the cache vault is re-loaded into memory, and the pending updates are submitted to the relevant logical units. In no case are partial I/O operations propagated to the logical unit.

Cache vault areas are themselves implemented in a RAID configuration. Thus, the vault area is protected against disk failures.

For DMX-2 arrays, battery backup power maintains power to all disk resources, and ensures that all I/O operations complete to disk. In extended power failure conditions, DMX-2 arrays will suspend access from host systems, destage current I/O operations to disk, and then power down.

### Mirrored write cache

For DMX-3 arrays, EMC implements 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.

### Cyclic redundancy checks

All update operations written to cache are fully protected by cyclic redundancy checks (CRC) within the array. This ensures protection against undetected faults within the array, such as intermittent data path errors. Data must pass CRC tests before read operations are successfully serviced by the array.

## Core 1.04 – Asynchronous capabilities

All EMC storage arrays under this compliance statement adhere to this requirement. EMC storage platforms will not transition asynchronous I/O operations from a host into synchronous operations.

## Core 1.05 – Write ordering

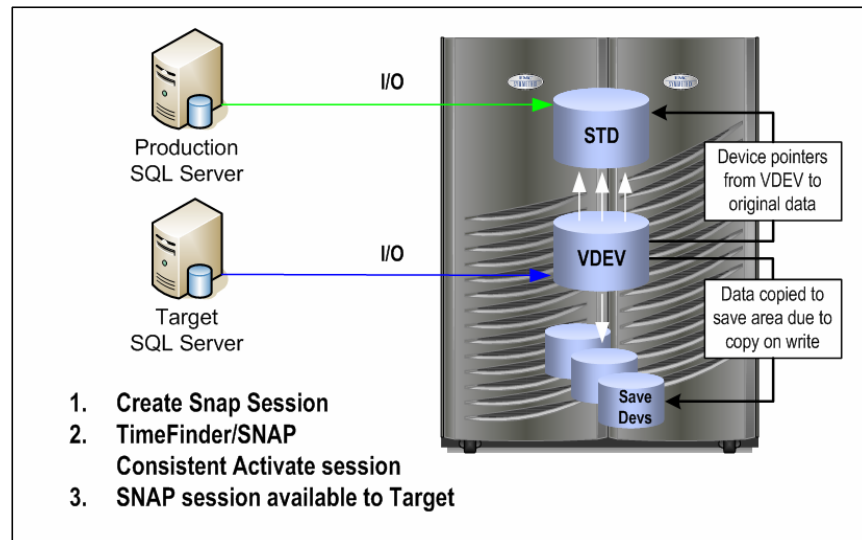
All EMC storage arrays under this compliance statement adhere to and can enforce write ordering.

### Non-replicated configurations

In non-replicated solutions, EMC Symmetrix storage arrays under this compliance statement honor the write dependency even in configurations where Microsoft SQL Server 2005 databases span storage arrays. Write ordering in this style of configuration is managed by SQL Server, and durability of I/O operations to stable media in each array is protected by compliance to “Core 1.02 – Stable media.”

Additionally, Symmetrix storage arrays provide support for EMC consistency technology to further extend protection of write order dependency. EMC consistency technology enables storage arrays to adhere to dependent write principles, which are the foundation of write ordering. Consistency groups can define related storage objects that need to be treated in an atomic manner to ensure that write ordering is protected. EMC Symmetrix arrays can implement consistency technology such that it internally maintains write ordering for operations such as EMC TimeFinder® consistent split operations. Figure 1 demonstrates the use of TimeFinder/SNAP technology using consistency technology within the same array. The logical object represented by both the “STD” and “VDEV” objects may actually comprise several storage devices (LUNs), which are the locations for the database data and log files. The resulting disk state presented to the target host will represent a restartable image, and complies with the Always On requirements.

In the same manner, all EMC TimeFinder implementations utilizing consistency technology would represent compliant solutions including those solutions utilizing TimeFinder/Mirror with business continuance volumes (BCVs) and TimeFinder/Clone.



**Figure 1 EMC TimeFinder/SNAP with consistency technology**

EMC Symmetrix storage arrays can provide consistency technology which spans multiple homogenous arrays. This provides adherence to “Core 1.05 – Write ordering” in situations where one or more SQL Server databases are defined across Symmetrix storage arrays.

Creation of consistency groups across heterogeneous array families is currently not supported.

## Replicated configurations

For Symmetrix storage solutions in replicated environments, the Symmetrix Remote Data Facility (SRDF) is a business continuity solution that provides a host-independent, mirrored data storage solution for duplicating production site data to one or more physically separated target Symmetrix systems. In basic terms, SRDF is a configuration of multiple Symmetrix units whose purpose is to maintain multiple copies of logical volume data in more than one location.

SRDF solutions that utilize consistency technology are fully compliant under the Microsoft SQL Server 2005 Always On specification.

SRDF replicates production or primary (source) site data to a secondary (target) site transparently to users, applications, databases, and host processors. The local SRDF device, known as the source (R1) device, is

configured in a partner relationship with a remote target (R2) device, forming an SRDF pair. While the R2 device is mirrored with the R1 device, the R2 device is write-disabled to the remote host. After the R2 device synchronizes with its R1 device, the R2 device can be split from the R1 device at any time, making the R2 device fully accessible again to its host. After the split, the target (R2) device contains valid data and is available for performing business continuity tasks through its original device address.

SRDF requires configuration of specific source Symmetrix volumes (R1) to be mirrored to target Symmetrix volumes (R2). If the primary site is no longer able to continue processing when SRDF is operating in synchronous mode, data at the secondary site is current up to the last I/O transaction. When primary systems are down, SRDF enables fast failover to the secondary copy of the data so that critical information becomes available in minutes. Business operations and related applications may resume full functionality with minimal interruption.

SRDF currently supports the following modes of operation:

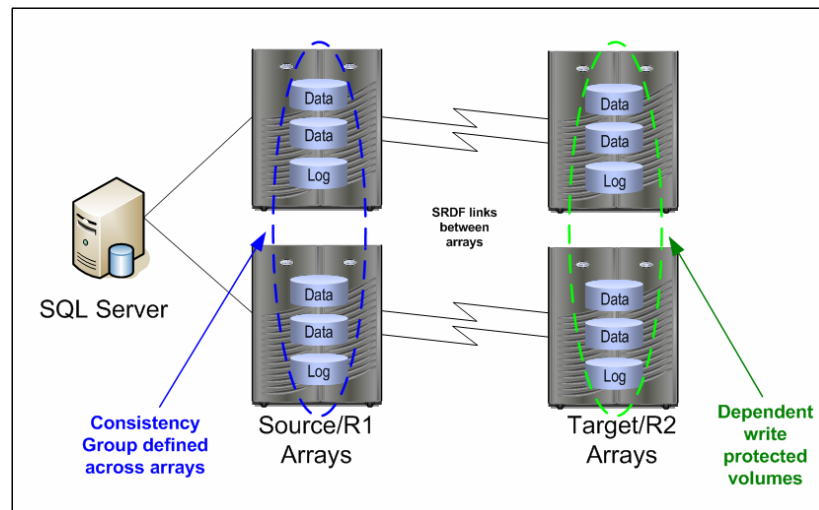
Synchronous mode (SRDF/S) provides real-time mirroring of data between the source Symmetrix system(s) and the target Symmetrix system(s). Data is written to the cache of both systems in real time before the application I/O is completed, ensuring the highest possible data availability. Data must be successfully stored in both the local and remote Symmetrix units before an acknowledgment is sent to the local host. This mode is used mainly for metropolitan area network distances less than 200 km.

Asynchronous mode (SRDF/A) maintains a dependent-write consistent copy of data at all times across any distance with no host application impact. Customers wanting to replicate data across long distances historically have had limited options. SRDF/A delivers high-performance, extended-distance replication and reduced telecommunication costs while leveraging existing management capabilities with virtually no host performance impact.

Adaptive mode transfers data from source devices to target devices regardless of order or consistency, and without host performance impact. This is especially useful when transferring large amounts of data during data center migrations, consolidations, and in data mobility environments.

SRDF adaptive copy mode does not support WAL and is only intended for use as a data migration facility.

To enable consistency within and across Symmetrix arrays, consistency groups are utilized. The purpose of the consistency group is to protect data integrity for applications that span multiple RA groups and/or multiple Symmetrix arrays. The protected applications may comprise multiple heterogeneous data resource managers across multiple host operating systems. Figure 2 details a logical view of a consistency group definition protecting write ordering across multiple arrays.



**Figure 2 Consistency group protection for database environment spanning arrays**

An SRDF consistency group uses PowerPath® or Enginuity Consistency Assist (ECA) support to provide synchronous disaster restart with zero data loss. Disaster restart solutions that use consistency groups provide remote restart with short recovery time objectives. Zero data loss implies that all completed transactions at the beginning of a disaster will be available at the target.

When the amount of data for an application becomes very large, the time and resources required for host-based software to protect, back up, or run decision-support queries on these databases become critical factors. The time required to quiesce or shut down the application for offline backup is no longer acceptable. SRDF consistency groups allow users to remotely mirror the largest data environments and automatically split off dependent-write consistent, restartable copies of applications in seconds without interruption to online service.

A consistency group is a composite group of SRDF devices (R1 or R2) that act in unison to maintain the integrity of applications distributed across multiple Symmetrix units or multiple RA groups

within a single Symmetrix. If a source (R1) device in the consistency group cannot propagate data to its corresponding target (R2) device, EMC software suspends data propagation from all R1 devices in the consistency group. This halts all data flow to the R2 targets. In the example provided in Figure 2, this suspension, called tripping the consistency group, would ensure that a dependent-write consistent R2 copy of the database up to the point in time that the consistency group tripped.

Tripping a consistency group can occur either automatically or manually. Scenarios in which an automatic trip would occur include:

- One or more R1 devices cannot propagate changes to their corresponding R2 devices

- The R2 device fails

- The SRDF directors on the R1 side or R2 side fail

In an automatic trip, the Symmetrix unit completes the write to the R1 device, but indicates that the write did not propagate to the R2 device. EMC software intercepts the I/O and instructs the Symmetrix to suspend all R1 source devices in the consistency group from propagating any further writes to the R2 side. Once the suspension is complete, writes to all of the R1 devices in the consistency group continue normally, but they are not propagated to the target side until normal SRDF mirroring resumes.

Symmetrix SRDF/A solutions enabled for use with consistency also represent valid solutions under the Microsoft SQL Server 2005 Always On environment. Figure 3 details the methodology for Symmetrix SRDF/A operations. As each switch between delta sets occurs in a dependent-write consistent manner, the resulting image is a valid restart point.

SRDF/A solutions can span multiple source storage arrays and represent valid solutions for databases spanning those arrays.

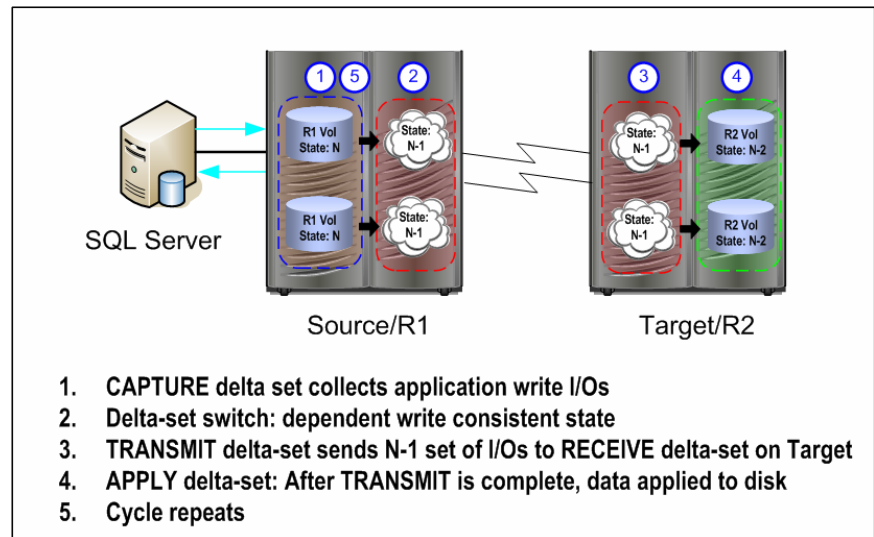


Figure 3 Symmetrix SRDF/A implementation utilizing dependant write consistency

### SRDF/CE for Microsoft Cluster Service

The EMC Geographically Dispersed cluster product, SRDF/CE for Microsoft Cluster Service (SRDF/CE for MSCS) utilizes SRDF/Synchronous connectivity for its implementation. SRDF/S as previously described follows the requirements for Microsoft SQL Server Always On compliance. Since all relevant technologies are implemented within the storage array, SRDF/CE for MSCS is a fully compliant configuration when Generic SafeWrite functionality is enabled as discussed in “Core 1.06 – Torn page protection.”

## Core 1.06 – Torn page protection

All EMC storage arrays under this compliance statement support the core requirement in the ways discussed in the remainder of this section.

Customers should also ensure that they implement best practice recommendations such as Partition Alignment as described within the *EMC Engineering Solutions Guide: Microsoft SQL Server on EMC Symmetrix Storage System*. Correct alignment will ensure that internal Microsoft SQL Server data structures are fully aligned to physical structures within the array. While not mandatory for compliance, alignment is highly recommended.

EMC Symmetrix storage solutions prevent torn page scenarios in configurations where the following Generic SafeWrite (GSW) functionality is implemented. The GSW functionality is utilized to

protect those volumes used by applications specifically from activities, which would result in torn page scenarios.

When storage volumes are protected with GSW functionality, the Symmetrix array will ensure that only full and complete I/O operations are accepted from the host and written to disk. In the event that a complete I/O operation is not received from the host, the Symmetrix will discard the I/O operation, and no acknowledgement will be sent to the host. The SCSI protocol requires that the host will re-try the I/O operation. Should the incomplete I/O operation have been caused by a failing operating system, the incomplete, and therefore invalid, I/O will never be written to disk.

A complete overview of the process and implementation of Generic SafeWrite is provided in the document *EMC Solutions Enabler Symmetrix Array Controls CLI Version 6.2* and later. A summary of the steps for configuring and then enabling the functionality is provided here.

### Configure devices for Generic SafeWrite

Before enabling Generic SafeWrite functionality, the RDB\_cksum Symmetrix device flag must be enabled on all devices targeted for Generic SafeWrite use. This change does not turn Generic SafeWrite on, it only allows it to be enabled on the specified devices.

The RDB\_cksum device flag can be set by using the SYMCLI symconfigure command, which will perform a Symmetrix configuration change.

The following is an example command:

```
symconfigure -sid 54 -f c:\enable_cksum.txt commit
```

where the c:\enable\_cksum.txt file contains the following command:

```
set device 0015:0019 attribute=RDB_Cksum;
```

### Enabling Generic SafeWrite

Once the device flags are set on the Symmetrix array, it is then possible to use the symchksum command to enable Generic SafeWrite.

Using the symchksum command, Generic SafeWrite can be enabled by specifying a specific device, a range of devices, or a device group.

To enable Generic SafeWrite for a device, use the command syntax shown in the example below:

```
symchksum enable -type generic dev 005 -sid 54
```

Note: If this is a metadvice, only the metahead needs to be specified.

To enable Generic SafeWrite for a contiguous range of devices the following syntax can be used:

```
symchksum enable -type generic -range 005:025 -sid 54
```

To assist in identifying the actual volumes to be configured with Generic SafeWrite, it is possible to utilize the Symmetrix Storage Resource Management (SRM) tools to identify database file locations and their corresponding Symmetrix devices. Additional information regarding the EMC SRM tools as they related to Microsoft SQL Server 2005 may be found in the *EMC Engineering Solutions Guide Microsoft SQL Server on EMC Symmetrix Storage Systems*. Refer to “Appendix B – References” for information on these additional resources.

## Core 1.07 – NTFS support

All EMC storage arrays covered under this compliance statement provide full support for all NTFS capabilities.

## Recommended 1.08 – Partner enhanced escalation

EMC is a Microsoft Gold Certified Support partner and both companies have signed a Cooperative Support Agreement (CSA). The agreement formally commits EMC 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 ver 1.02

### Requirements

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

Type	Definition
Required	A capability or property that the subsystem <i>must</i> provide to pass the requirements of the Always On Storage Solution Review Program.
Recommended	A capability or property that 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
<p>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.</p> <p>For detailed information, see the "Power Outage Testing – Pull The Plug" section in <i>Microsoft SQL Server I/O Basics</i>, Chapter 2.</p>	

Core 1.03 Forced Unit Access (FUA) and Write-Through	Required
<p>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.</p> <p>All components in a solution must honor the write-to-stable media intent. This includes, but is not limited to, caching components.</p> <p>It is not enough to only honor WAL for Microsoft SQL Server log files. Data files and backup streams also depend on WAL behavior.</p> <p>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.</p> <p>For more information, see 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.</p>	

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 about how a synchronous action can affect the Microsoft SQL Server scheduler, see the white paper, <i>How To Diagnosis and Correct Errors 17883, 17884, 17887, and 17888</i>.</p>	

Core 1.05 Write Ordering	Required
<p>A tenant 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 the 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 that requires 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><b>Example Configuration</b></p> <p>A solution has the following configuration:</p> <ul style="list-style-type: none"> <li>• Data File—Device A—Subsystem #1</li> <li>• Log File—Device B—Subsystem #2</li> </ul> <p>If these subsystems use separate physical paths with different caching, Microsoft SQL Server would not be able 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, see 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, we 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, see the “Torn I/O,” “Log Parity,” and “Sector Size” sections located in the white paper, <i>Microsoft SQL Server I/O Basics</i>, Chapter 2.</p>	

Core 1.07 NTFS Support	Required
<p>You must support NTFS capabilities. This includes but not limited to the following:</p> <ul style="list-style-type: none"> <li>• Sparse Files</li> <li>• File Streams</li> <li>• Encryption</li> <li>• Compression</li> <li>• All Security Properties</li> </ul> <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 1.07 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* (applies to SQL Server 7.0, SQL Server 2000, and SQL Server 2005)

### EMC customer documentation

EMC provides documentation to existing customers through the EMC Powerlink™ site (<http://Powerlink.EMC.com/>). The following documentation is available in the documentation library:

For information on Symmetrix Storage Arrays and integration with Microsoft SQL Server:

*EMC Engineering Solutions Guide: Microsoft SQL Server on EMC Symmetrix Storage Systems*

For information regarding implementing Generic SafeWrite:

*EMC Solutions Enabler: Symmetrix Array Controls CLI Version 6.2*