Windows Server 2012 Certified

HITACHI
Inspire the Next

TECHNICAL PAPER

Microsoft SQL Server I/O Reliability Storage System on Hitachi Virtual Storage Platform G600

Certification Report

By Steven Burns

August 2015

Hitachi Data Systems

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Contents

Microsoft SQL Server I/O Reliability Storage System on Hitachi Virtual Storage Platform G600

Certification Report

This document describes how Hitachi Virtual Storage Platform G600 complies with the Microsoft® SQL Server® I/O Reliability Program (SSRP) for Microsoft SQL Server 2014. Each requirement has been addressed for compliance and a detailed statement made for each. The SQLIOSim and TPC-E tests were conducted at the Hitachi Data Systems labs with the configurations required from Microsoft to ensure full compliance with all the requirements as described in SQL Server I/O Reliability Program Review Requirements for SQL Server 2014.

The SQL Server I/O Reliability Program reviews storage systems against core functional SQL Server storage requirements. Storage systems that meet SSRP requirements are reliable and highly available for SQL Server deployment. The program does not define or review the performance characteristics of storage solutions, but does require two I/O tests using the following:

- SQLIOSim
- TPC-E Test Kit

This report provides information about requirements and compliance based on an update to the SSRP program not available online when publishing this document. These requirements were provided to Hitachi Data Systems by the SQL Server team from Microsoft.

Program Requirements and Compliance

These are the requirements from Microsoft for a storage system to be accepted as a Microsoft SQL Server I/O Reliability Program member and how Hitachi Virtual Storage Platform G600 complies with those requirements for SQL Server 2014.

Core 1.00: Windows Logo Certification

This describes the core Microsoft Windows logo certification requirements.

Requirements

Microsoft Windows logo certification helps ensure the safety of Microsoft SQL Server data by testing various aspects of the Windows logo program. To be compliant with the SQL Server I/O Reliability Solution Review Program, solutions must pass and maintain the latest certifications for Windows logos.

The Windows logo program has hardware and software tracks. SQL Server I/O Reliability requires completion and logo certification for the tests applicable to the I/O solution.

Compliance

Hitachi Data Systems has a well-established process that ensures every microcode and storage system release earns the Windows logo program. Contact Hitachi Data Systems or visit Microsoft Windows Server Catalog to review these reports. On **Windows Server Catalog**, search for Hitachi Data Systems.

Core 1.01: Core Windows API Support

This describes the core Microsoft Windows API support requirements.

Requirements

Microsoft SQL Server utilizes several APIs to enable 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. Table 1, starting on page 3, outlines the required storage APIs and the fundamental properties of each API, which must be supported. Table 2 on page 5 outlines the Core 1.01 API behavioral needs requirements.

Table 1. Core 1.01 API Fundamental Requirements

Core API	Fundamentals
CreateFile	SQL Server uses many of the options allowed by CreateFile for various operations on database, backup, security, and other files. The following are specific options of importance. ■ FILE_FLAG_WRITETHROUGH ■ FILE_FLAG_NO_BUFFERING ■ FILE_FLAG_OVERLAPPED ■ FILE_FLAG_RANDOM_ACCESS ■ FILE_FILE_SEQUENTIAL_SCAN ■ FILE_SHARE_READ ■ FILE_SHARE_WRITE ■ GENERIC_READ ■ GENERIC_WRITE ■ READ_CONTROL ■ WRITE_DAC ■ WRITE_OWNER ■ SECURITY_ANONYMOUS ■ SECURITY_SQOS_PRESENT ■ SECURITY_VALID_SQOS_FLAGS ■ SERVER_SQOS_MANDATORY Recommended: Support for secondary stream usage. Secondary stream usage and naming is outlined in SQL Server I/O Basics, Chapter 2. SQL Server 2014 extends online DBCC behavior for ReFS.

Table 1. Core 1.01 API Fundamental Requirements (Continued)

Core API	Fundamentals
DeviceIoControl	DeviceIoControl is used to establish and query file system behaviors. Discovering details about volume information, establishing sparse file settings, obtaining sparse file allocation information, or other Attributes. The following are specific options and data structures of importance. The I/O system should ensure accuracy of the returned data. ■ FSCTL_SET_SPARSE ■ FSCTL_SET_ZERO_DATA ■ FSCTL_QUERY_ALLOCATED_RANGES ■ FSCTL_SET_COMPRESSION ■ IOCTL_DISK_GET_LENGTH_INFO ■ IOCTL_DISK_GET_PARTITION_INFO_EX ■ IOCTL_DISK_GET_DRIVE_GEOMETRY ■ GPT_BASIC_DATA_ATTRIBUTE_SHADOW_COPY ■ GPT_BASIC_DATA_ATTRIBUTE_READ_ONLY Note: DeviceIoControl can be called by using and without using the OVERLAPPED structure.
FlushFileBuffers GetVolumePathName GetVolumeInformation GetVolumeNameForVolumeMountPoint	The solution must honor the SL_WRITE_THROUGH and IRP_MJ_FLUSH_BUFFERS. ■ Used to harden a backup before the transaction log is truncated. ■ Used when file size changes are made to ensure metadata is flushed. ■ Enables support of mount points. ■ Provides details about the volume, similar to those outlined in the DeviceIoControl entry.

Table 2. Core 1.01 API Behavioral Needs Requirements

Core API	Behavioral Needs
WriteFile	■ OVERLAPPED Present ■ OVERLAPPED Not-Present ■ General I/O size ranges from 512-bytes to 4MB
WriteFileGather	■ OVERLAPPED Present ■ OVERLAPPED Not-Present ■ General I/O size ranges from 8KB to 8MB ■ Systems that do not support scatter/gather can experience performance degradation. ■ Used heavily by SQL Server lazy writer, eager writes (bulk operations), and checkpoint processing ■ WOW64 on X64 support needed
ReadFile	■ OVERLAPPED Present ■ OVERLAPPED Not-Present ■ General I/O size ranges from 512-bytes to 4MB
ReadFileScatter	■ OVERLAPPED Present ■ OVERLAPPED Not-Present ■ General I/O size ranges from 8KB to 10MB ■ Systems that do not support scatter/gather can experience performance degradation ■ Used heavily by SQL Server read ahead logic ■ WOW64 on X64 Support needed

SQLIOSim is a utility for use by storage solution providers to provide a means of testing. A storage solution provider may use SQLIOSim to help review and test an implementation. See “Core 1.08: Testing” on page 10 for testing requirements.

The NTFS file-system has two ways to tell the subsystem to flush the data to stable media:

- Marking each write IRP as SL_WRITE_THROUGH
- Sending IRP_MJ_FLUSH_BUFFERS

When FlushFileBuffers is called, NTFS will write each page remaining in file system cache and invoke IRP_MJ_FLUSH_BUFFERS. The flush must ensure all data has been written to stable media.

Just using FLAG_WRITE_THROUGH + FILE_FLAG_NO_BUFFERRING does not cause IRP_MJ_FLUSH_BUFFERS, although SL_WRITE_THROUGH is set on every IRP.

Compliance

Hitachi Data Systems compliance for this section is shown in the SQLIOSim and TPC-E Kit profiles provided by Microsoft. This requirement is described in more detail in Core 1.08: Testing.

Hitachi Virtual Storage Platform G600 complies with these requirements because they can be accessed using block-based iSCSI or Fibre Channel protocols. This includes full support for the following:

- Overlapped I/O
- Sparse files
- Direct and buffered or delayed I/O calls

Core 1.02: Stable Media

This describes the core stable media requirements.

Requirements

SQL Server relies on the Write-Ahead Logging (WAL) protocol to do the following:

- Maintain the ACID properties of the database:
  - Atomicity
  - Consistency
  - Isolation
  - Durability

WAL relies on stable media capabilities. A solution must comply with this stable media intention. For detailed information, see the “Pull-the-plug power outage testing” in SQL Server I/O Basics, Chapter 2.

Compliance

All writes to Hitachi Virtual Storage Platform G600 are committed to battery-backed, nonvolatile mirrored cache and then committed to RAID protected volumes within the storage system.

Hitachi Virtual Storage Platform G600 is protected from data loss by means of dual power supplies for each controller. Each controller has a cache that is backed by a battery. If a power failure continues more than 20 milliseconds, the storage system uses power from the batteries to back up the cache memory data and the storage system configuration data onto the cache flash memory.

Core 1.03: Forced Unit Access (FUA) and Write-Through

This describes the core forced unit access (FUA) and write-through requirements.

Requirements

To support Write-Ahead Logging (WAL), Microsoft SQL Server uses FILE_FLAG_WRITETHROUGH when opening database files. SQL Server also uses FlushFileBuffers during various operations. Both of these options must be Write-through and flushing to stable media 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.

SQL Server log files must honor WAL. In addition, data files and backup streams also depend on WAL behavior.

Many storage products include battery-backed caching mechanisms. If these caching mechanisms are present in the solution, the SQL Server I/O Reliability solution white paper should document the practical limits of the write-through stable media protection for a production environment.

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

Compliance

Hitachi Virtual Storage Platform G600 complies with the FUA and Write-through requirements.

Hitachi Virtual Storage Platform G600 use of the FILE_FLAG_WRITETHROUGH flag ensures that data is written directly from SQL Server in a non-buffered fashion to stable media immediately. In the case that data is buffered on the host, a batch operation will use the FlushFileBuffers call to flush the data to stable media. In the event of a failure in the storage system during the flush process, the SQL Server host is notified of a failure to write backup data to disk.

The architecture of Hitachi Virtual Storage Platform G600 ensures that all write data within the battery-backed cache subsystem is preserved as well as redundant. To accomplish this, the following functions are implemented:

- **Mirrored Cache for Writes** — With all Hitachi Virtual Storage Platform G600, each cache has its own data region and a mirrored data region of the other cache in the event of a failure. Not only are all data writes stored in the mirrored region, but all reads are as well.
- **Battery Backup** — Hitachi Virtual Storage Platform G600 has a battery backup system designed to ensure that all cache write data is preserved and that all I/O is properly de-staged to the disk storage system.
- **Monitoring** — Hi-Track® Remote Monitoring system from Hitachi Data Systems provides automated monitoring of the batteries. This “call home” service and remote maintenance tool provides 24/7 notifications to IT service personnel in the event of a battery failure.

Core 1.04: Asynchronous Capabilities

This describes the core asynchronous capability requirements.

Requirements

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 SQL Server I/O Reliability solution must provide asynchronous I/O capabilities.

For more information about how a synchronous action can affect the Microsoft SQL Server scheduler, see How to Diagnosis and Correct Errors 17883, 17884, 17887, and 17888.

Compliance

Hitachi Virtual Storage Platform G600 complies with the asynchronous capabilities requirements in “Core 1.04: Asynchronous Capabilities ” on page 7. It complies with the requirements needed to support NTFS including write ordering. This is explained in more detail in “Core 1.05: Write Ordering ” on page 8 and “Core 1.07: NTFS Support ” on page 9.

Core 1.05: Write Ordering

This describes the core write ordering requirements.

Requirements

A tenet of the WAL protocol is write-ordering preservation. Any SQL Server I/O Reliability solution must provide write ordering preservation.

For more information about write ordering requirements, see “Write Ordering” and “Mirroring and Remote Mirroring” in SQL Server 2000 I/O Basics.

Compliance

As mentioned in “Core 1.02: Stable Media” on page 6, Hitachi Virtual Storage Platform G600 complies with the WAL protocol requirements including the write ordering preservation. SQL Server 2000 I/O Basics has the expanded definition of this requirement that SQL Server I/O Basics Chapter 2 references. The compliance review was done against the expanded definition set forth in SQL Server 2000 I/O Basics.

All writes issued to Hitachi Virtual Storage Platform G600 are written in the order they are received. In the case of replication technologies, all data remains consistent from the originating source data. Further explanation of this is in “Advanced 2.01: Write Ordering — Required for Remote Storage Solutions” on page 11.

This article was written originally for Microsoft SQL Server 2000, but the I/O basics remain the same on newer versions of SQL Server.

Core 1.06: Torn I/O Protection

This describes the core torn I/O protection requirements.

Requirements

A SQL Server I/O Reliability solution must provide sector alignment and sizing in a way that prevents torn I/O, including splitting I/Os across various I/O entities in the I/O path.

Additionally, a solution must accurately report sector size to Windows I/O APIs. Accurately reporting sector size helps prevent sector size mismatches and 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 or exposed as torn writes. Any SQL Server I/O Reliability solution must document configurations in such a way that use actual sector sizes from the sector size list that is supported by Microsoft SQL Server: 512, 1024, 2048, and 4096 bytes.

To indicate when a torn-write situation occurs, Microsoft recommends that the solution log appropriate warning events.

The SQL Server I/O Reliability solution white paper must include information about the configuration requirements needed for the solution to meet the torn I/O requirements.

For more information, see SQL Server I/O Basics, Chapter 2.

SSD/Flash (Non-Spinning Media)

Sector sizes become opaque (often simulated) to the I/O stack for solutions which are not based on spinning media. Solutions with non-spinning media designs must still avoid torn writes. These solutions should document how torn writes are avoided. They should also document how to configure the solution to reduce Read/Modify/Write I/O patterns, reduce wear of the media and optimize I/O performance.

Compliance

Hitachi Virtual Storage Platform G600 fully supports standard NTFS sector sizing, torn I/O requirements and sector sizing reporting requirements (512 bytes, 1024 bytes, 2048 bytes, and 4096 bytes). For Microsoft Windows Server 2008 release and later, sector alignment is not necessary due to the use of 1024K as the default partition offset. However for Microsoft Windows Server 2003, use Diskpart.exe to ensure that the disk tracks are sector and cache aligned.

The required testing defined in section Core 1.08: Testing was performed on SAS hard drives and on SSDs to demonstrate compliance for both drive types.

Core 1.07: NTFS Support

This describes the NTFS requirements.

Requirements

You must support NTFS capabilities. This includes but is not limited to the following:

- Sparse Files
- File Streams
- Encryption
- Compression
- All Security Properties

The solution must support sparse files on NTFS based file systems. Microsoft SQL Server 2005 and newer versions uses sparse files in support of DBCC CHECK* commands and snapshot databases.

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. Storage solution providers may choose to provide utilities to copy or move sparse files without destroying the sparse file intent.

Compliance

Hitachi Virtual Storage Platform G600 complies with all the NTFS support requirements described in this section, which also closely align with the Windows Logo Certification program. Hitachi Virtual Storage Platform G600 provides full support for these features using Fibre Channel and iSCSI protocols.

Core 1.08: Testing

This describes the core testing requirements.

Requirements

The SQL Server I/O Reliability Review Program requires successful execution of the following data durability test suites:

- SQLIOSim
  - The latest version of SQLIOSim was used with the configuration file provided by Microsoft.
  - Stress testing ran for 24 hours.
- TPC-E Test Kit

  The test kit ran with the following configuration as required by Microsoft:
  - Database protection level (PAGE AUDIT) set to checksum.
  - Trace flags (818, 815) were enabled.

The following criteria must be met on both tests:

- Low paged and non-paged pool conditions
- Excessive outstanding request boundaries. For example 10000+.
- Memory requests are not forced to a specific memory location. For example, a 64 bit driver requiring a memory allocation of less than 4 GB can force aggressive working set trimming.

The following optional conditions can be reviewed for additional level of validation:

- x86 with 3GB Enabled (optional)
- x86 with PAE Enabled (optional)
- WOW64 running x86 in x64 (optional)

Compliance

Hitachi Data Systems performed the tests required to demonstrate compliance with these requirements on a Hitachi Virtual Storage Platform G600 storage system. This testing applies to all midrange systems of the Hitachi Virtual Storage Platform family because they all use the same microcode and maintain the same level of compliance with the requirements described in this and other sections.

Hitachi Data Systems used the SQLIOSim profile provided by Microsoft to execute the 24-hour stress test on the storage system. The required testing generated no errors or warnings.

The benchmark suite test required for this section was executed with the use of the TPC-E benchmark kit provided by Microsoft. The tests execution completed successfully and generated no warnings or errors.

Advanced 2.01: Write Ordering — Required for Remote Storage Solutions

This describes the advanced write ordering requirements.

Requirements

For remote and mirrored I/O destinations, all the paths must honor write ordering across the database files. A SQL Server I/O Reliability 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.”

Note — If these subsystems use separate physical paths with different caching, Microsoft SQL Server may not be able to support this configuration because the caching mechanisms may not present a coherent cache. The subsystems may require a third element to maintain cache coherency across the disparate caches. Consider including system databases in the consistency group to provide enhanced metadata consistency.

Caching problems can also occur across network boundaries. If a database backup is written to a UNC path but FlushFileBuffers only ensures that the local system file cache is flushed, SQL Server may be exposed to data loss. Network based solutions must ensure stable media delivery.

For non-battery backed cache solutions, the following must be met:

- If the solution uses cache that is not battery backed, it must provide stable media and write ordering guarantees as well. An exception is TEMPDB.
- This type of solution often holds the write buffers in RAM memory during transmission and must use a two phase commit approach to maintain stable media and write ordering requirements.

Compliance

Hitachi Virtual Storage Platform G600 provides high availability and disaster recovery features that comply with the remote storage capabilities requirements from Microsoft. Hitachi offers the following software products for high availability and disaster recovery on Hitachi Virtual Storage Platform G600:

- Hitachi TrueCopy®
  - TrueCopy modules offer a group-based update sequence consistency solution. This enables fast and accurate database recovery, even after a “rolling” disaster, without the need for time-consuming data recovery procedures.
  - The synchronous capabilities provide a volume based real-time data backup and are ideal for high-priority data backup, duplication, and migration tasks. In the event of a disaster or system failure at the primary site, the secondary synchronous data can be invoked to allow recovery at the volume level with a high level of data integrity.
  - The synchronous capabilities write data to cache in the primary storage system and the secondary storage system prior to signaling the application that the I/O is complete. TrueCopy volume groups at the remote site can be recovered with full update sequence consistency.
  - The synchronous capabilities of TrueCopy software modules provide update sequence consistency for user defined groups of volumes, such as databases and associated logs, as well as protection for write-dependent applications in the event of a failure.

- Hitachi ShadowImage® Heterogeneous Replication provides in system, full volume clones with consistency across user defined groups of volumes for local or remote point in time protection.
- Hitachi Thin Image snapshot software provides in system, differential based copies with consistency across user defined groups of volumes for local or remote point in time protection.
- Hitachi Universal Replicator provides asynchronous, write order consistent remote replication with user defined consistency groups.

These Hitachi, Ltd., and Hitachi Data Systems software products are critical parts of disaster recovery solutions for large amounts of data that span multiple volumes and storage systems across long distances.

Advanced 2.02: Transactional Sector/Block Rewrites

This describes the advanced transactional sector or block rewrites requirements.

Requirements

Solutions involving movement of sectors or blocks must provide transactional safety while maintaining asynchronous capabilities. Sectors or blocks cannot be rewritten or changed unless transactional safety can be guaranteed.

See SQL Server 2000 I/O Basics for details.

Compliance

Hitachi Virtual Storage Platform G600 complies with these requirements, which are a subset of these requirements, particularly regarding retention of a database's ACID properties:

- "Core 1.02: Stable Media" on page 6
- "Core 1.03: Forced Unit Access (FUA) and Write-Through" on page 6
- "Core 1.04: Asynchronous Capabilities ," starting on page 7

In addition, Hitachi Virtual Storage Platform G600 has a pool optimization function that provides a zero page reclamation option that is part of Hitachi Dynamic Provisioning. When a user reclaims zero data during the optimization of a given dynamic provisioning pool, a page on the Hitachi Virtual Storage Platform G600 consisting solely of binary zeroes is unmapped. If enough pages across all the chunks owned by a dynamic provisioning volume (DP-VOL) are unmapped, the optimization process recombines allocated pages, thereby releasing a chunk and increasing the free capacity of the dynamic provisioning pool.

Hitachi Tiered Storage Manager uses Hitachi TrueCopy asynchronous capabilities that enable the efficient use of storage tiers through either automatic (scripted) or manual execution. The Hitachi TrueCopy asynchronous replication used in Hitachi Tiered Storage Manager ensures sectors and blocks are replicated from primary to secondary volumes while ensuring transactional safety. Data can be moved between tiers for performance, archiving, or data migration needs.

The page or LU is copied from source to target but the key is that the two sides must be in a paired state before the active data pointer is changed from the source to the target volume during the tier migration process. This ensures that the data and the copy process have transactional safety.

On Hitachi Virtual Storage Platform G600, TrueCopy synchronous replication and Hitachi Universal Replicator can be used by Hitachi Tiered Storage Manager to create a duplicate production volume on a secondary volume. Data in a TrueCopy backup stays synchronized with the primary data. This happens when data is written from the host to the local storage system then to the remote system using a Fibre Channel or iSCSI link. The host holds subsequent output until it receives an acknowledgment from the secondary storage system that a copy of the previous output was completed.

Advanced 2.03: VDI

This describes the VDI requirements.

Requirements

VDI solutions must meet the requirements outlined in the SQL Server 2005 Virtual Backup Device Interface (VDI) Specification.

Compliance

Hitachi complies with these requirements with Hitachi Replication Manager using VDI through the component object model (COM). Figure 1 illustrates how Hitachi Replication Manager works with VDI in a Microsoft SQL Server environment.

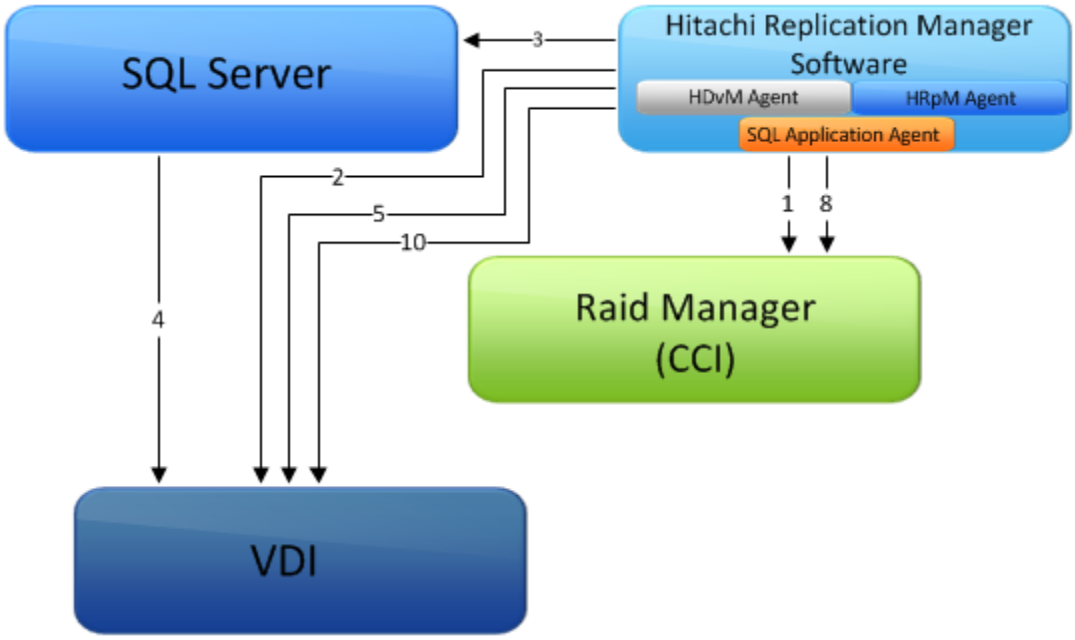

Figure 1

In the diagram:

- HRpM stands for Hitachi Replication Manager
- HDvM stands for Hitachi Device Manager
- CCI stands for command control interface
- VDI stands for Microsoft SQL Server Virtual Device Interface

The following steps show how Hitachi Replication Manager works with VDI in a SQL Server environment:

1. Replication Manager resynchronizes the target copy pairs.
2. Replication Manager opens VDI through the COM.
3. Replication Manager runs an SQL Server statement specifying VDI.
4. Microsoft SQL Server sends backup definition to VDI.
5. Replication Manager receives the backup definition from VDI and extracts the definition to the metafile.
6. Microsoft SQL Server freezes the write I/O of the target database.
7. Replication Manager purges cache buffer (FlushFileBuffers()).
8. Replication Manager splits the target copy pairs.
9. Microsoft SQL Server thaws the I/O.
10. Replication Manager closes VDI.

Hitachi Data Systems also offers Hitachi Storage Adapter for Microsoft Volume ShadowCopy Service, a provider for Microsoft Volume Shadow Copy Service (VSS), which complies with the VSS requirements from Microsoft and works with Microsoft SQL Server. VSS is part of Microsoft File and Storage services framework and enables the backup infrastructure for the Microsoft Windows operating system and applications. VSS provides the ability to perform I/O consistent, point-in-time backups using storage assisted technologies, like Hitachi ShadowImage Replication.

Hitachi Storage Adapter for Microsoft Volume ShadowCopy Service integrates the in-system replication products with the Microsoft VSS framework to create the following:

- Clones using ShadowImage Replication
- Snapshots using Hitachi Thin Image for Windows-based applications, such as Microsoft Exchange and SQL Server.

Storage Adapter for Microsoft Volume ShadowCopy Service also splits the pair during the backup operation so an off-host backup can take place.

Advanced 2.04: Clustering

Requirement deleted.

Advanced 2.05: File Streams

This describes the file streams requirements.

Requirements

The Microsoft SQL Server file streams feature requires NTFS transactional guarantees. Compliance of SQL Server file streams requires the following:

- File system must report NTFS.
- File system must support the NTFS mini-filter stack including support of the filter manager contexts FSRTL_FLAG2_SUPPORTS_FILTER_CONTEXTS in FCBs.
- File stream access does not support OpLocks
- Support for Extended Attributes.

Metadata changes on the same volume must maintain ordering. For example, a metadata change of File A followed by a metadata change of File B will maintain order, even after a crash recovery.

Directory scan enumeration returns the current state of the directory. For example, if a directory contains File A and File C when the scan is started but, during the scan, file B is added; then File A and File C should be returned and File B is optional. The file system must not look at the count of files at scan startup as an absolute.

In the example, the count would be two (2). The scan should not return File A and File B only. The enumeration behavior of File B is undefined but File A and File C were present at the start of the scan and unaltered during the scan. So File A and File C should be returned in the scan.

It is recommended that the solution participates in Microsoft Plug Fest Interop testing.

Compliance

Hitachi Virtual Storage Platform G600 complies with these requirements. This includes complying with all the NTFS support requirements described in this section, which closely align with the Windows Server Logo Certification program.

Hitachi does not offer any proprietary file system to their customers that would mandate compliance with this requirement.

Advanced 2.06: Protection

This describes the protection requirements.

Requirements

Data durability compromises can frequently be predicted or avoided. This is frequently referred under initiatives such as S.M.A.R.T. Microsoft recommends that solutions provide advanced data protection features.

Custom stability checks can be implemented in conjunction with an Always On solution.

Compliance

Hitachi Virtual Storage Platform G600 uses hard drives that provide advanced data protection features established by the S.M.A.R.T. initiative. To use S.M.A.R.T., a system OEM writes software that reads the S.M.A.R.T. status at the ATA or SCSI interface and presents the S.M.A.R.T. status to the end user.

In the latest S.M.A.R.T. revision, commands can do the following in compliance with the requirements of the S.M.A.R.T initiative:

- Perform some self-testing of the hard disks used in Hitachi Virtual Storage Platform G600
- Implement S.M.A.R.T technology in the hard drive microcode

Advanced 2.07: Hardware Virtualization

This describes the hardware virtualization requirements.

Requirements

Solutions involving virtualized environments must comply with Server Virtualization Validation Program, (SVVP) as outlined in Support policy for Microsoft SQL Server products that are running in a hardware virtualization environment.

Compliance

Hitachi logical partitioning complies with the SVVP program requirements. It is listed on the vendor list available on Server Vendor.

Conclusion

This report shows that the Hitachi Virtual Storage Platform G600 is in full compliance with all of the requirements defined and set forth in Microsoft SQL Server I/O Reliability Program (SSRP) for Microsoft SQL Server 2014. This is supported with the successful completion and review of the results of the SQLIOSim and TPC-E Kit tests by Microsoft.

References

These are the references used in this document.

- Microsoft KB articles:
  - http://support.microsoft.com/kb/917043
  - http://support.microsoft.com/kb/917047
  - http://support.microsoft.com/kb/956893
- Microsoft SQL Server
  - http://www.microsoft.com/en-us/sqlserver/default.aspx
- MSDN (Microsoft Developer Network) Blogs
  - http://blogs.msdn.com/
- Microsoft Download Center
  - http://www.microsoft.com/downloads
- Microsoft Server and Cloud Platform
  - http://www.microsoft.com/en-us/server-cloud/windows-server/networking.aspx
- Hitachi Data Systems
  - http://www.hds.com/assets/pdf/hitachi-architecture-guide-hus-vm.pdf
  - http://www.hds.com/products/storage-software/shadowimage-in-system-replication.html
  - http://www.hds.com/products/storage-software/universal-replicator.html
- Wikipedia
  - http://en.wikipedia.org/wiki/S.M.A.R.T.

For More Information

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Hitachi Data Systems Global Services website.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate Resources website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services Education website.

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the Hitachi Data Systems website.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 96050-2639 USA
www.HDS.com community.HDS.com

Regional Contact Information
Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

AS-408-00, August 2015