

[MS-SQLRS-Diff]:

SQL Server Remote Storage Profile

Intellectual Property Rights Notice for Open Specifications Documentation

- **Technical Documentation.** Microsoft publishes Open Specifications documentation (“this documentation”) for protocols, file formats, data portability, computer languages, and standards support. Additionally, overview documents cover inter-protocol relationships and interactions.
- **Copyrights.** This documentation is covered by Microsoft copyrights. Regardless of any other terms that are contained in the terms of use for the Microsoft website that hosts this documentation, you can make copies of it in order to develop implementations of the technologies that are described in this documentation and can distribute portions of it in your implementations that use these technologies or in your documentation as necessary to properly document the implementation. You can also distribute in your implementation, with or without modification, any schemas, IDLs, or code samples that are included in the documentation. This permission also applies to any documents that are referenced in the Open Specifications documentation.
- **No Trade Secrets.** Microsoft does not claim any trade secret rights in this documentation.
- **Patents.** Microsoft has patents that might cover your implementations of the technologies described in the Open Specifications documentation. Neither this notice nor Microsoft's delivery of this documentation grants any licenses under those patents or any other Microsoft patents. However, a given Open Specifications document might be covered by the Microsoft [Open Specifications Promise](#) or the [Microsoft Community Promise](#). If you would prefer a written license, or if the technologies described in this documentation are not covered by the Open Specifications Promise or Community Promise, as applicable, patent licenses are available by contacting iplg@microsoft.com.
- **License Programs.** To see all of the protocols in scope under a specific license program and the associated patents, visit the [Patent Map](#).
- **Trademarks.** The names of companies and products contained in this documentation might be covered by trademarks or similar intellectual property rights. This notice does not grant any licenses under those rights. For a list of Microsoft trademarks, visit www.microsoft.com/trademarks.
- **Fictitious Names.** The example companies, organizations, products, domain names, email addresses, logos, people, places, and events that are depicted in this documentation are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

Reservation of Rights. All other rights are reserved, and this notice does not grant any rights other than as specifically described above, whether by implication, estoppel, or otherwise.

Tools. The Open Specifications documentation does not require the use of Microsoft programming tools or programming environments in order for you to develop an implementation. If you have access to Microsoft programming tools and environments, you are free to take advantage of them. Certain Open Specifications documents are intended for use in conjunction with publicly available standards specifications and network programming art and, as such, assume that the reader either is familiar with the aforementioned material or has immediate access to it.

Support. For questions and support, please contact dochelp@microsoft.com.

Revision Summary

Date	Revision History	Revision Class	Comments
8/16/2017	1.0	New	Released new document.
9/15/2017	2.0	Major	Significantly changed the technical content.
3/16/2018	2.0	None	No changes to the meaning, language, or formatting of the technical content.
10/16/2019	3.0	Major	Significantly changed the technical content.
11/1/2022	4.0	Major	Significantly changed the technical content.

Table of Contents

1	Introduction	4
1.1	Glossary	4
1.2	References	5
1.2.1	(Updated Section) Normative References.....	5
1.2.2	(Updated Section) Informative References.....	5
1.3	(Updated Section) Microsoft Implementations.....	5
1.4	Standards Support Requirements	5
1.5	Notation.....	5
2	Standards Support Statements.....	7
2.1	Normative Variations	7
2.2	Clarifications	7
2.2.1	[MS-SMB2] Server Message Block (SMB) Protocols Version 2 and 3	7
2.2.2	[MS-FSCC] File System Control Codes and [MS-FSA] File System Algorithms	7
2.2.3	[MS-FSCC] File System Control Codes.....	8
2.3	Error Handling	8
2.4	Security	8
3	Change Tracking.....	9
4	Index.....	10

1 Introduction

The SQL Server Remote Storage Profile specifies information regarding the implementation for hosting Microsoft SQL Server databases on Server Message Block (SMB) shares. The files associated with the storage of databases can include data files, transaction log files, and other associated on-disk artifacts. This profile clarifies the level of support that SQL Server requires from SMB servers that host these types of files.

Section 2 in this specification is normative. All other sections and examples in this specification are informative.

1.1 Glossary

This document uses the following terms:

file system: A system that enables applications to store and retrieve files on storage devices. Files are placed in a hierarchical structure. The file system specifies naming conventions for files and the format for specifying the path to a file in the tree structure. Each file system consists of one or more drivers and DLLs that define the data formats and features of the file system. File systems can exist on the following storage devices: diskettes, hard disks, jukeboxes, removable optical disks, and tape backup units.

file system control (FSCTL): A command issued to a file system to alter or query the behavior of the file system and/or set or query metadata that is associated with a particular file or with the file system itself.

I/O control (IOCTL): A command that is issued to a target file system or target device in order to query or alter the behavior of the target; or to query or alter the data and attributes that are associated with the target or the objects that are exposed by the target.

Resilient File System (ReFS): The Resilient File System is a file system that provides maximum data availability, scalability, and data integrity despite hardware or software failures. ReFS is frequently used together with Storage Spaces.

shadow copy: A duplicate of data held on a volume at a well-defined instant in time.

snapshot: The point in time at which a shadow copy of a volume is made.

sparse file: A file containing large sections of data composed only of zeros. This file is marked as a sparse file in the file system, which saves disk space by only allocating as many ranges on disk as are required to completely reconstruct the non-zero data. When an attempt is made to read in the nonallocated portions of the file (also known as holes), the file system automatically returns zeros to the caller.

stream: A sequence of bytes written to a file on the target file system. Every file stored on a volume that uses the file system contains at least one stream, which is normally used to store the primary contents of the file. Additional streams within the file can be used to store file attributes, application parameters, or other information specific to that file. Every file has a default data stream, which is unnamed by default. That data stream, and any other data stream associated with a file, can optionally be named.

volume: A group of one or more partitions that forms a logical region of storage and the basis for a file system. A volume is an area on a storage device that is managed by the file system as a discrete logical storage unit. A partition contains at least one volume, and a volume can exist on one or more partitions.

MAY, SHOULD, MUST, SHOULD NOT, MUST NOT: These terms (in all caps) are used as defined in [RFC2119]. All statements of optional behavior use either MAY, SHOULD, or SHOULD NOT.

1.2 References

Links to a document in the Microsoft Open Specifications library point to the correct section in the most recently published version of the referenced document. However, because individual documents in the library are not updated at the same time, the section numbers in the documents may not match. You can confirm the correct section numbering by checking the Errata.

1.2.1 (Updated Section) Normative References

We conduct frequent surveys of the normative references to assure their continued availability. If you have any issue with finding a normative reference, please contact dohelp@microsoft.com. We will assist you in finding the relevant information.

[MS-FSA] Microsoft Corporation, "File System Algorithms".

[MS-FSCC] Microsoft Corporation, "File System Control Codes".

[MS-SMB2] Microsoft Corporation, "Server Message Block (SMB) Protocol Versions 2 and 3".

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <https://www.rfc-editor.org/rfc/rfc2119.txt>

1.2.2 (Updated Section) Informative References

[MSDN-IOCTLSQPCC] Microsoft Corporation, "IOCTL_STORAGE_QUERY_PROPERTY IOCTL", https://docs.microsoft.com/en-us/windows/desktop/api/winiocpl/ni-winiocpl-ioctl_storage_query_property

[MSDN-IOCTLVGGACC] Microsoft Corporation, "IOCTL_VOLUME_GET_GPT_ATTRIBUTES IOCTL", https://docs.microsoft.com/en-us/windows/desktop/api/winiocpl/ni-winiocpl-ioctl_volume_get_gpt_attributes

1.3 (Updated Section) Microsoft Implementations

- Microsoft SQL Server 2016
- Microsoft SQL Server 2017
- Microsoft SQL Server 2019
- Microsoft SQL Server 2022

1.4 Standards Support Requirements

The conformance requirements for [MS-SMB2], [MS-FSCC], and [MS-FSA] are that all required portions of the specifications are implemented according to the specification, and that any optional portions that are included are also implemented according to the specification.

1.5 Notation

The following notations are used to identify clarifications in section 2.2.

Notation	Explanation
C####	This notation identifies a clarification of ambiguity in the target specification. This includes imprecise statements, omitted information, discrepancies, and

Notation	Explanation
	errata. This does not include data formatting clarifications.

2 Standards Support Statements

2.1 Normative Variations

None.

2.2 Clarifications

The following subsections identify clarifications relative to [MS-SMB2], [MS-FSCC], and [MS-FSA].

2.2.1 [MS-SMB2] Server Message Block (SMB) Protocols Version 2 and 3

C0001:

The server MUST support the Server Message Block (SMB) version 3.0 or higher dialect, as described in [MS-SMB2] section 1.7.

C0002:

The server MUST support persistent handles.

2.2.2 [MS-FSCC] File System Control Codes and [MS-FSA] File System Algorithms

C0003:

If the server supports the FSCTL_FILESYSTEM_GET_STATISTICS control code, as defined in [MS-FSCC] section 2.3.12 and [MS-FSA] section 2.1.5.9.7, and if the server responds with a FILESYSTEM_STATISTICS structure, as defined in [MS-FSCC] section 2.3.12.1, with a FileSystemType field value of "FILESYSTEM_STATISTICS_TYPE_REFS", SQL Server issues a query on the stream to retrieve file system integrity information, as defined in [MS-FSCC] section 2.3.19 and [MS-FSA] section 2.1.5.9.10.

C0004:

If the server supports the FSCTL_SET_INTEGRITY_INFORMATION control code, as defined in [MS-FSCC] section 2.3.73 and [MS-FSA] section 2.1.5.9.33, and if the underlying remote file system is a Resilient File System (ReFS), support for this file system control (FSCTL) code can improve the performance of remote input/output (I/O) operations. This control code is used by SQL Server to identify and disable ReFS integrity streams.

C0005:

If the server supports the FSCTL_GET_INTEGRITY_INFORMATION control code, as defined in [MS-FSCC] section 2.3.19 and [MS-FSA] section 2.1.5.9.10, and if the underlying file system is a ReFS file system, support for this control code can improve the performance of remote I/O operations. This control code is used by SQL Server to identify and disable ReFS integrity streams.

C0006:

If the server supports the FSCTL_SET_SPARSE control code, as defined in [MS-FSCC] section 2.3.83 and [MS-FSA] section 2.1.5.9.38, and if the SMB remote share is used to host SQL Server data files, core SQL Server engine features, such as database snapshots and online database consistency checks, use sparse files to reduce the amount of disk space that is required to store copy-on-write snapshots of SQL Server data.

If the remote file system does not support this control code, database snapshots and online database consistency checks fail.

2.2.3 [MS-FSCC] File System Control Codes

C0007:

If the server supports the `IOCTL_VOLUME_GET_GPT_ATTRIBUTES` control code, as defined in [MS-FSCC] section 2.8.3, SQL Server can query properties of the volume that are underlying the remote file share. This query can occur even if the volume is a shadow copy, such as a volume snapshot, or is read-only. For further information, see [MSDN-IOCTLVGGACC].

If the server does not support this I/O control (IOCTL) code, SQL Server is unable to detect these properties and some operations could fail or behave unexpectedly.

C0008:

If the server supports the `IOCTL_STORAGE_QUERY_PROPERTY` control code, as defined in [MS-FSCC] section 2.8.1, with the **StorageAccessAlignmentProperty** property ID, SQL Server can identify the physical sector size of any device that is underlying the SQL Server database and transaction log files and, by extension, the required alignment of SQL Server I/O requests. For further information, see [MSDN-IOCTLSQPCC].

If the server does not support this control code, SQL Server cannot correctly identify the physical sector size and could suffer performance degradations due to read-modify-write operations.

2.3 Error Handling

None.

2.4 Security

None.

3 Change Tracking

This section identifies changes that were made to this document since the last release. Changes are classified as Major, Minor, or None.

The revision class **Major** means that the technical content in the document was significantly revised. Major changes affect protocol interoperability or implementation. Examples of major changes are:

- A document revision that incorporates changes to interoperability requirements.
- A document revision that captures changes to protocol functionality.

The revision class **Minor** means that the meaning of the technical content was clarified. Minor changes do not affect protocol interoperability or implementation. Examples of minor changes are updates to clarify ambiguity at the sentence, paragraph, or table level.

The revision class **None** means that no new technical changes were introduced. Minor editorial and formatting changes may have been made, but the relevant technical content is identical to the last released version.

The changes made to this document are listed in the following table. For more information, please contact dochelp@microsoft.com.

Section	Description	Revision class
1.3 Microsoft Implementations	Added SQL Server 2022 to the list of Microsoft implementations.	Major

4 Index

C

Change tracking 9
Clarifications - standards support statements 7

E

Error handling – standards support statements 8

G

Glossary 4

I

Informative references 5
Introduction 4

M

MS-FSCC - standards support statements 8
MS-FSCC and MS-FSA - standards support statements 7
MS-SMB2 - standards support statements 7

N

Normative references 5
Normative variations - standards support statements 7

R

References
 informative 5
 normative 5

S

Security – standards support statements 8
Standards support statements
 clarifications 7
 error handling 8
 MS-FSCC 8
 MS-FSCC and MS-FSA 7
 MS-SMB2 7
 normative variations 7
 security 8

T

Tracking changes 9