



APPENDIX **A**

SCSI Write Acceleration and Tape Acceleration

This appendix describes the concepts of SCSI write acceleration, tape acceleration, and compression.

This appendix includes the following sections:

- [SCSI Write Acceleration, page A-1](#)
- [SCSI Tape Acceleration, page A-2](#)


SCSI Write Acceleration

The SCSI write acceleration feature enables you to significantly improve application write performance when storage traffic is routed over wide area networks using FCIP or Fibre Channel. When write acceleration is enabled, WAN throughput is maximized by minimizing the impact of WAN latency for write operations.

In [Figure A-1](#), the WRITE command without write acceleration requires two round-trip transfers (RTT), while the WRITE command with write acceleration only requires one RTT. The maximum sized Transfer Ready is sent from the host side of the FCIP or Fibre Channel link back to the host before the WRITE command reaches the target. This enables the host to start sending the write data without waiting for the long latency over the FCIP or Fibre Channel link of the WRITE command and Transfer Ready. It also eliminates the delay caused by multiple Transfer Readys needed for the exchange going over the FCIP or Fibre Channel link.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

Figure A-1 SCSI Write Acceleration



105224

SCSI Tape Acceleration

Tapes are storage devices that store and retrieve user data sequentially. Cisco MDS NX-OS provides both tape write and read acceleration.

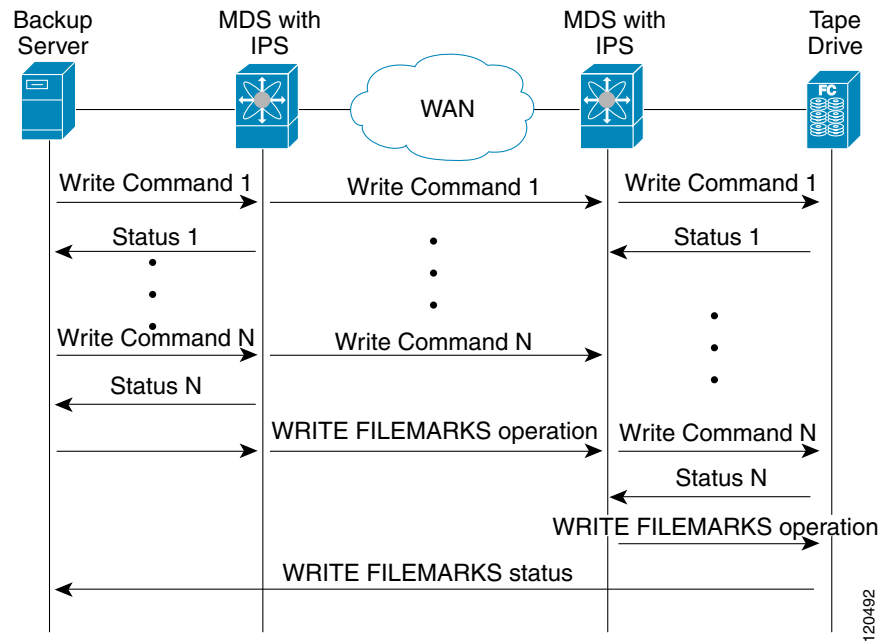
Applications that access tape drives normally have only one SCSI WRITE or READ operation outstanding to it. This single command process limits the benefit of the tape acceleration feature when using an FCIP or FC tunnel over a long-distance WAN link. It impacts backup, restore, and restore performance because each SCSI WRITE or READ operation does not complete until the host receives a good status response from the tape drive. The SCSI tape acceleration feature helps solve this problem. It improves tape backup, archive, and restore operations by allowing faster data streaming between the host and tape drive over the WAN link.

In an example of tape acceleration for write operations, the backup server in [Figure A-2](#) issues write operations to a drive in the tape library. Acting as a proxy for the remote tape drives, the local Cisco MDS switch proxies a transfer ready to signal the host to start sending data. After receiving all the data, the local Cisco MDS switch proxies the successful completion of the SCSI WRITE operation. This

Send documentation comments to mdsfeedback-doc@cisco.com

response allows the host to start the next SCSI WRITE operation. This proxy method results in more data being sent over the FCIP or Fibre Channel tunnel in the same time period compared to the time taken to send data without proxying. The proxy method improves the performance on WAN links.

Figure A-2 SCSI Tape Acceleration for Write Operations



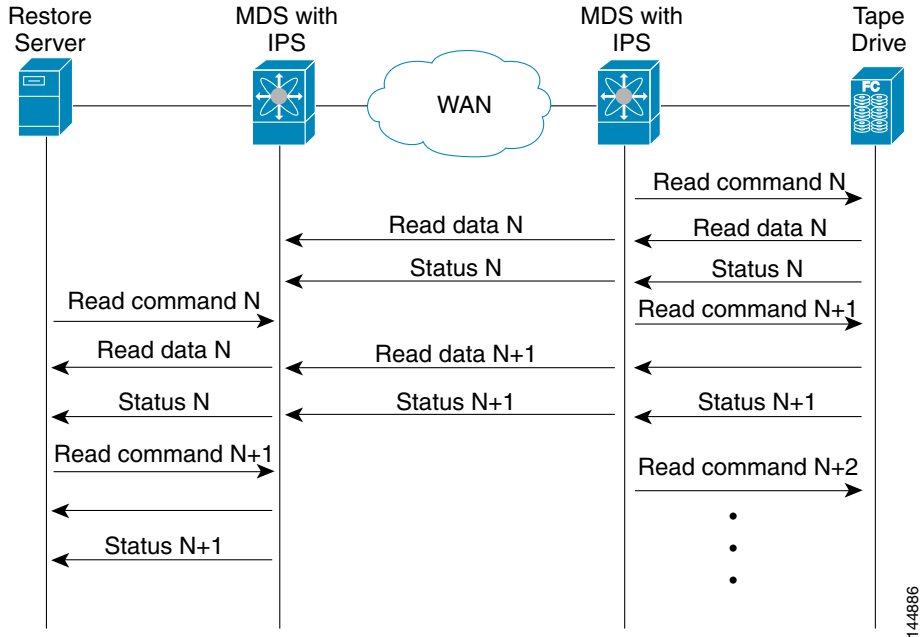
At the tape end of the FCIP or Fibre Channel tunnel, another Cisco MDS switch buffers the command and data it has received. It then acts as a backup server to the tape drive by listening to a transfer ready from the tape drive before forwarding the data.

The Cisco NX-OS provides reliable data delivery across the WAN. It maintains write data integrity by allowing the WRITE FILEMARKS operation to complete end-to-end without proxying. The WRITE FILEMARKS operation signals the synchronization of the buffer data with the tape library data. While tape media errors are returned to backup servers for error handling, tape busy errors are retried automatically by the Cisco NX-OS software.

In an example of tape acceleration for read operations, the restore server in [Figure A-3](#) issues read operations to a drive in the tape library. During the restore process, the remote Cisco MDS switch at the tape end, in anticipation of more SCSI read operations from the host, sends out SCSI read operations on its own to the tape drive. The prefetched read data is cached at the local Cisco MDS switch. The local Cisco MDS switch on receiving SCSI read operations from the host, sends out the cached data. This method results in more data being sent over the FCIP or FC tunnel in the same time period compared to the time taken to send data without read acceleration for tapes. This improves the performance for tape reads on WAN links.

Send documentation comments to mdsfeedback-doc@cisco.com

Figure A-3 SCSI Tape Acceleration for Read Operations



The Cisco NX-OS provides reliable data delivery across the WAN. While tape media errors during the read operation are returned to the restore server for error handling, the Cisco NX-OS software recovers from any other errors.

In tape acceleration for writes, after a certain amount of data has been buffered at the remote Cisco MDS switch, the write operations from the host are flow controlled by the local Cisco MDS switch by not proxying the Transfer Ready. On completion of a write operation when some data buffers are freed, the local Cisco MDS switch resumes the proxying. Likewise, in tape acceleration for reads, after a certain amount of data has been buffered at the local Cisco MDS switch, the read operations to the tape drive are flow controlled by the remote Cisco MDS switch by not issuing any further reads. On completion of a read operation, when some data buffers are freed, the remote Cisco MDS switch resumes issuing reads.

The default flow control buffering uses the **automatic** option. This option takes the WAN latencies and the speed of the tape into account to provide optimum performance. You can also specify a flow control buffer size (the maximum buffer size is 12 MB).