

# In Brief

## Optimizing Storage Data Transfers over Long Distance WAN Links with Fast Write

### INTRODUCTION

SCSI is a simple high-speed protocol for transferring large amounts of storage data. SCSI was originally designed with the assumption that the SCSI client and SCSI server are in close proximity of each other. However, today SCSI is used over Fibre Channel and Fibre Channel over IP, to transfer storage data over distances that vary from a few feet to thousands of miles. When SCSI is used to transport data over distances that introduce latencies of the order of milliseconds, the performance of SCSI write operations can be significantly impacted. McDATA's Fast Write capability is designed to address the performance issues of SCSI Write operations over long distance, high latency links.

### THE BASIC SCSI WRITE OPERATION

SCSI read operations are rather simple; the SCSI client initiates a request for data and the SCSI server responds with all of the requested data without any further acknowledgements or handshakes. However, the SCSI write operation is more involved in that it may require multiple handshakes between the SCSI client and SCSI server. Before a SCSI client sends any data to the SCSI server, the client has to find out how much buffer space is available on the server. The SCSI client sends a message to the SCSI server indicating the size of the data it wants to write. The SCSI server then responds with a ready-to-transfer (FC\_XFER\_RDY) message specifying how much data the client is allowed to transfer, without further acknowledgement. Once the client sends the negotiated amount of data, the SCSI server responds with another FC\_XFER\_RDY, if there is more data left to transfer as part of that Write operation. The SCSI client and server continue to interact with data and FC\_XFER\_RDY messages until all of the data for the entire write operation has been transferred from the client to server. Only at that point does the server send the command completion message back to the client, acknowledging that it has received and stored all of the information written by the client. This process is illustrated in Figure 1.

### DISTANCE AND LATENCY AFFECT SCSI WRITE OPERATION PERFORMANCE

When the SCSI write operation is performed over distance, each additional round-trip message handshake between the SCSI client and server increases the time needed to complete the write operation. A greater number of message handshakes require a longer length of time to complete the write operation. In addition, performance is dependent on the distance and latency involved in the long distance data transfer, since the greater latency adds time to each roundtrip handshake required to complete the write operation. Figure 3 illustrates the impact of distance and latency on performance for long distance write operations. As a result of the distance and latency effects, long distance SCSI write operations can suffer significant performance degradation. Unless measures are taken to mitigate the impacts of latency, storage data transfers for distances of beyond 100 km will suffer to a degree considered unacceptable for a large number of storage applications.

# In Brief

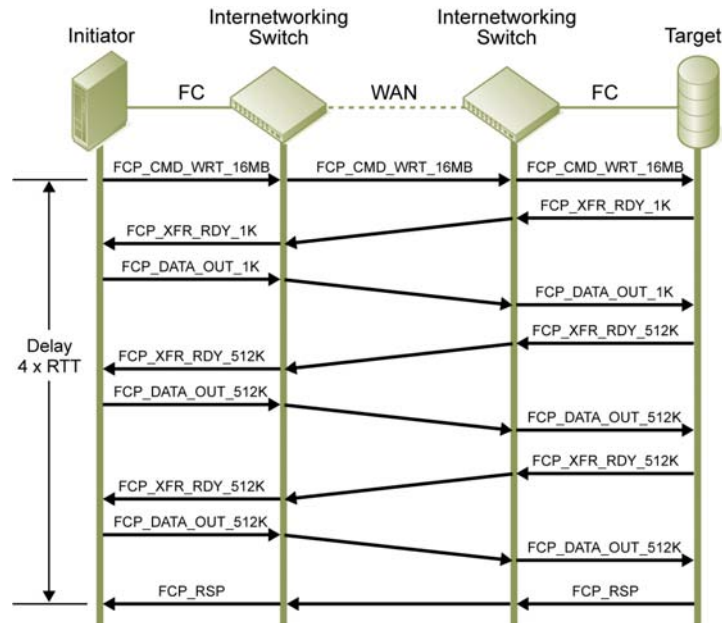
## Optimizing Storage Data Transfers over Long Distance WAN Links with Fast Write

### THE MCDATA FAST WRITE FEATURE

The McDATA Fast Write feature is designed to overcome the latency effects for SCSI write operations, without any compromise to data integrity and security. Fast Write allows the entire data segment of the SCSI operation to be transported across the long distance link without the inefficiencies of waiting for the ready-to-transfer (FC\_XFER\_RDY) acknowledgements to travel back and forth across the high-latency environment.

### HOW FAST WRITE WORKS

Figure 2 illustrates the mechanics of McDATA's Fast Write feature. When a write operation is detected, the McDATA IPS switch forwards the write command to the target normally. Commands are therefore delivered to the target in the same order that they were issued by the initiator. However, the McDATA IPS switch, acting as a virtual target, immediately issues a ready-to-transfer (FC\_XFER\_RDY) message to the initiator, prompting it to transmit the entire data segment for the write operation. The MCDATA IPS switches then transfer the data across the high latency environment to the remote McDATA IPS switch. The remote target device then interacts with the remote McDATA IPS switch, which acts as a virtual initiator. Ready-to-transfer messages issued by the target are handled directly by the remote McDATA IPS switch, as if it were the real initiator issuing the data for the write operation.



*Figure 1: SCSI write operations often involve multiple handshake messages between target and initiator to transfer a SCSI write data segment from the initiator to the target. When long distances and significant latencies exist between target and initiator, performance can be adversely impacted.*

# In Brief

## Optimizing Storage Data Transfers over Long Distance WAN Links with Fast Write

### NO DANGER OF DATA CORRUPTION

There is no possibility of data corruption, since all data is protected by the CRC in the Fibre Channel frames. All commands are received by the target in the same order that they were issued by the initiator. Furthermore, Fast Write does not interfere with the final command completion message issued by the target device. Any error or "check condition" status issued by the target device would be detected through delivery (or non-delivery) of the final SCSI operation completion message.

### SIGNIFICANTLY ENHANCED PERFORMANCE

Fast Write has been demonstrated to provide anywhere between two-fold to an excess of a tenfold increase in aggregate SCSI performance, depending on a variety of factors including latency of the link, size of the SCSI write command and number of outstanding I/Os. Fast Write benefits are particularly significant for asynchronous applications such as mirroring and remote copy, which issue large SCSI I/O blocks of data.

### CONCLUSION

Fast Write is an exclusive capability available in IPS switches from McDATA Corporation that provide superior, potentially line-rate gigabit performance over extended distances. Fast Write allows the user to maximize utilization over expensive WAN links, without introducing additional risk to data integrity and security.

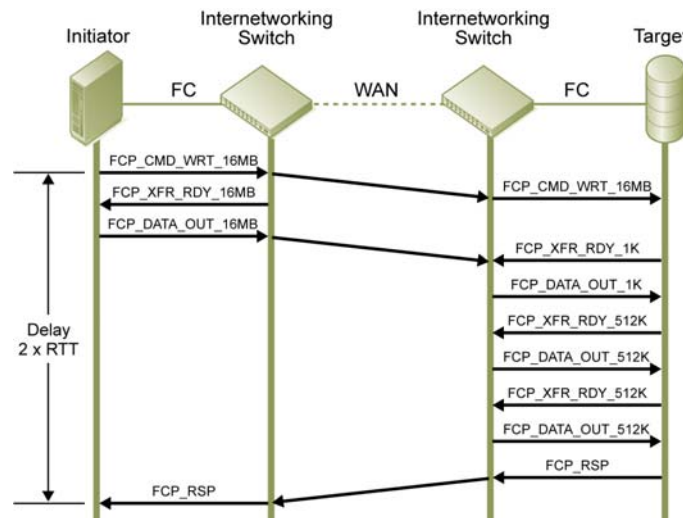


Figure 2: The Fast Write capability allows the McDATA IPS switches to expedite transfer of the SCSI write data segment, without having to wait for potentially numerous round-trip handshake messages to travel back and forth between target and initiator.

For more information about McDATA Corporation, please visit our web site at [www.mcdata.com](http://www.mcdata.com).