BROCADE

SAN TECH NOTE

Optimizing Storage Data Transfers over FC ISLs and Long-Distance FCIP WAN Links with FastWrite

INTRODUCTION

SCSI is a simple, high-speed protocol for transferring large amounts of storage data. It was designed on the assumption that the SCSI application client and SCSI device server would be in close proximity to each other. Currently, however, SCSI is used over Fibre Channel (FC) and Fibre Channel over Internet Protocol (FCIP) 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. The Brocade FastWrite feature addresses the performance issues of SCSI write operations over long-distance, high-latency links.

THE BASIC SCSI WRITE OPERATION

SCSI read operations are simple: the SCSI initiator initiates a request for data and the SCSI target responds with all of the requested data without further acknowledgements or round-trips across the connection. However, the SCSI write operation is more involved—it requires two round-trips between the SCSI initiator and target. Before a SCSI initiator is permitted to send any data to the SCSI target, the initiator must notify the target of the impending write to ascertain if space is available in the receiving buffer. The target responds with a "transfer ready" message if space is available. With most modern storage subsystems, this is merely a formality due to large cache sizes and efficient cache management algorithms. When a SCSI initiator intends to write an Information Unit (IU) it sends a message, which is part of an exchange called a sequence, to the SCSI target indicating the size of the write. The SCSI target then responds with a "transfer ready" (FCP_XFER_RDY) sequence specifying how much data the initiator is allowed to transfer, which is usually the size of the entire write. After all the data has been transferred, the target sends the command completion sequence (FCP_RSP) back to the initiator, acknowledging that it has received and stored all of the information written by the initiator. Upon receipt of that response, the write is complete.

When long distances and significant latencies exist between target and initiator, SCSI write operations can involve multiple handshake messages between target and initiator to transfer a SCSI write data sequence (FCP_DATA) from the initiator to the target, as shown in Figure 1.

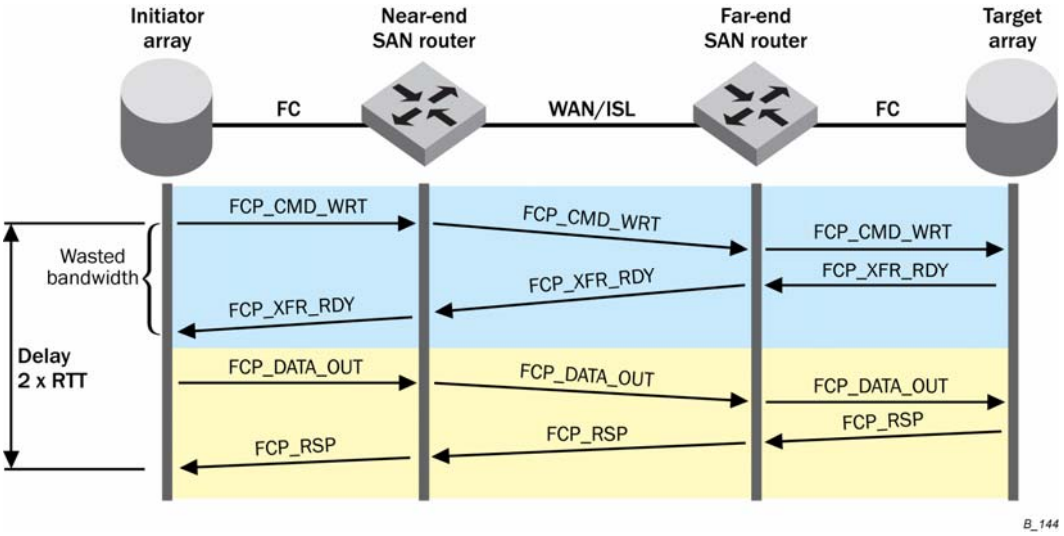

Figure 1. Performance can be adversely impacted with long-distance, high-latency data transfers.

Distance and Latency Impact SCSI Write Performance

When the SCSI write operation is performed over distance, each additional round-trip communication between the SCSI application client and device server increases the time needed to complete the overall write operation. This results in degraded application performance. Unless measures are taken to mitigate the negative impact of latency, storage data transfers over a few kilometers in distance will suffer to a degree considered unacceptable for a large number of storage applications.

BROCADE FASTWRITE

The Brocade FastWrite feature is designed to overcome the latency effects for SCSI write operations, without compromising data integrity and security. FastWrite allows the entire data sequence of the SCSI operation to be transported across the link, without the inefficiencies of waiting for the "transfer ready" (FCP_XFER_RDY) to travel back across the high-latency environment. Brocade FastWrite is available with either FC-based or FCIP extension.

Figure 2 illustrates the mechanics of Brocade FastWrite. When a write operation is detected, the Brocade 7500 SAN Router forwards the write command to the target in the standard way. Commands are therefore delivered to the target in the same order that they were issued by the initiator. However, the Brocade 7500, acting as a virtual target, immediately issues an FCP_XFER_RDY sequence to the initiator, prompting it to transmit the entire data sequence (FCP_DATA) for the write operation. The Brocade 7500 transfers the data across the high-latency environment to the remote SAN router. The remote target device then interacts with the remote Brocade SAN Router, which acts as a virtual initiator.

FCP_XFER_RDY issued by the target is handled directly by the remote Brocade SAN Router, as if the router were the real initiator issuing the data for the write operation. FastWrite allows the Brocade SAN Routers to expedite transfer of the SCSI write data sequence, without having to wait for potentially numerous round-trip handshake messages to travel back and forth between target and initiator.

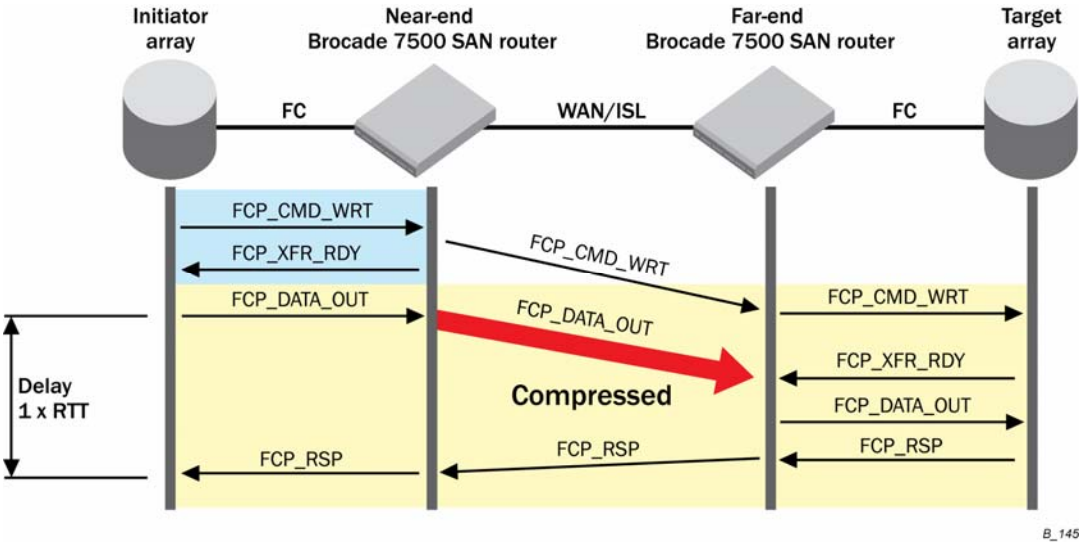

Figure 2. Brocade SAN Routers use FastWrite to expedite transfer of the SCSI write data segment.

No Danger of Data Corruption

There is no possibility of undetected data corruption, since all data is protected by a Cyclic Redundancy Check (CRC) in each Fibre Channel frame. All commands are received by the target in the same order that they were issued by the initiator, guaranteed by the nature of FC switching or TCP in the case of FCIP. Furthermore, FastWrite 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 is detected through delivery (or non-delivery) of the final SCSI operation completion message (FCP_RSP).

Significantly Enhanced Performance

FastWrite has been demonstrated to provide from a twofold to a greater than tenfold increase in aggregate SCSI performance, depending on factors such as latency of the link, size of the SCSI write command, and the number of outstanding I/Os. FastWrite benefits are significant for both synchronous and asynchronous applications, such as mirroring and remote copy. Asynchronous applications that cannot supply enough I/O to fill the connection benefit from data transfer starting in half the time. Synchronous applications benefit from a response time reduced by half.

Conclusion

FastWrite is a Brocade capability available in Brocade SAN Routers, which provides superior, line-rate gigabit performance over extended distances. FastWrite allows users to maximize utilization over FC Inter-Switch Links (ISLs) as well as expensive Wide Area Network (WAN) links, without introducing additional risk to data integrity and security, even over distances of just a few kilometers. FastWrite provides benefit for both native FC and FCIP by either increasing throughput or reducing response time by half.

For more information about Brocade SAN solutions, visit www.brocade.com.