

FIBRE CHANNEL NEW TECHNOLOGIES: FC-NVME-2

By Rupin Mohan, FCIA Marketing Chairman; Director R&D and CTO, Hewlett-Packard Enterprise

Fibre Channel is a serial point-to-point protocol that transfers a ton of data super reliably using light over fibre cables. Data is transferred from one port to another from the initiator port on the server to the destination port on the storage array.

A switching fabric allows communication among more than just initiator and target ports by introducing switching ports in the data path. Even if switch ports are in the data path, Fibre Channel is still a point-to-point protocol. The data transfers are reliable because just like an airplane only takes off after it has confirmed a landing spot at the destination airport, Fibre Channel data packets are only transmitted after a destination port confirms it has enough data buffers to receive the packet.

Fibre Channel does a great job of delivering data reliably. It doesn't matter if the data packets are

SCSI or NVMe. NVMe over Fibre Channel (FC-NVMe) defines the mapping of NVMe over Fabrics (NVMe-oF) to the Fibre Channel protocol.

The I/O operation defined by FC-NVMe is mapped into a Fibre Channel Exchange. A Fibre Channel Exchange carrying information for an NVMe Express over Fabrics I/O operation is an FC-NVMe Exchange. In simple English, the initiator ports sitting inside servers and the destination ports on the storage arrays establish their point-to-point connections and transfer data back and forth using Fibre Channel exchanges.

In order to provide a more reliable and solid base platform for NVMe over Fabrics, FC-NVMe-2 will introduce optimizations on how Fibre Channel Exchanges carry NVMe traffic. These optimizations include:

Exchanges are given some more time before getting terminated

FC-NVMe associations are terminated only after transmitting proper ABTS (Abort Sequence) so that all resources associated with this exchange are recovered before the exchange or the association is completely terminated.

Exchange association identifiers are reused more efficiently after exchanges are terminated

By properly terminating an outstanding exchange, as mentioned above, all the resources associated with the terminated exchange can be reused efficiently

Enhancements when supported by both initiator and target ports including:

Support first burst, further optimizing FC-NVMe data transfers

- Initiator may choose to perform faster write operations by sending first data without waiting for a transfer ready (XFER_RDY) message from the target. This helps in faster data transfers as long as the first burst data is within the limits of the agreed upon burst size.

Sequence Level Error Recovery (SLER), which allows higher level exchange recovery

- Sequence-level error recovery allows error recovery by using sequence re-transmissions. This is accomplished via some new commands (FLUSH and Responder Error Detected - RED) exchanged between initiator and target ports to indicate, and recover from, any sequence level errors during an exchange.

Support confirmed completion

- FC-NVMe-2 will define a mechanism so that the target and initiator ports can use special messages to determine successful completion of all sequences within a given exchange.

End-to-end data protection

Exchanges are retried when there is an error in data transfer

- As mentioned above, confirmed completion features when used together with SLER will result in better error handling and recovery.

Fibre Channel exchanges and sequences are the foundation of the protocol and with these new features, FC has built even a stronger foundation with even more reliability for the next generation storage protocol, NVMe over Fabrics. With all these enhancements, it's difficult to dispute that Fibre Channel is the most reliable and high-performance fabric protocol for NVMe packets.