

Fibre Channel Delivers So Much More Than Speed

By Marc Angelinovich, Principal Product Marketing Manager, Broadcom Inc.

For decades now, Fibre Channel has been the network of choice for storage when deploying critical applications like ERP to run highly complex and large organizations or financial applications at the largest banks and global stock exchanges. Reliability is the main reason why Fibre Channel has won over the competition time after time. Anyone can deliver performance numbers in a test environment, but when customers are running all resources and maximum performance and scale, reliability is tested. This brings me to the reason for this article, why beyond the need for speed, a Fibre Channel network is required for any enterprise that relies on their storage infrastructure to conduct critical business operations like revenue-generating applications.

Regardless of the vendor of choice, Fibre Channel is designed with the customer in mind first. Even though Ethernet providers would like to say the same thing, they can't. It starts at the foundation of why Fibre Channel exists. Fibre Channel is a purpose-built network with the sole purpose to ensure storage traffic connectivity regardless of distance, performance degradation, physical issues or any other issues that come up. Ethernet networks are designed for many different things like connecting manufacturing equipment, servers, smart devices and the list goes on, making it hard for a general-purpose network as an industry collective to focus on storage.

Fibre Channel delivers so much more than speed to the largest companies in the world. If the design goal of the network is to ensure storage traffic in the data center is always available, reliable and secure, then more often than not, Fibre Channel is the network of choice.

So, let's talk about the areas the Fibre Channel industry focuses on beyond speed. First is fabric services that any device can subscribe to. The second is end-to-end communication between devices. The third is reliability across the industry. All of these capabilities enables the Fibre Channel industry to do so much more like share actionable intelligence across all devices in the data center to ensure the greatest reliability.

Fabric services

Fabric services are a set of functions shared across different Fibre Channel vendors to provide centralized capabilities to build the foundation for simple discovery, access control (security), and management. Fabric services are provided by the following, standardized well-known servers and controllers:

- **Fabric Controller** – Facilitates the exchange of information between switches in the fabric.
- **Name Server (Directory)** – Provides a means to discover information about the end device, host bus adapter (HBA) and the port attached to a Fabric.
- **Management Server** – Provides a single management access point within the fabric for services such as fabric configuration, access control management (e.g. zoning and unzone name server), security policy distribution, device management and application services).
- **Domain Controller** – Facilitates features and functions unique to each switch.

Fabric services mean, if you have a Marvell or Broadcom Host Bus Adapter (HBA), a Cisco or a Brocade switch, and any Fibre Channel storage array, these devices will log in and register with each other and share common capabilities. Saying it a different way, each device will call out if they support functions like hardware signaling if congestion is identified or if they can support the identification of virtual machines. Once each device understands each other's capabilities, they can start working together.

These shared Fabric services are the foundation for customers to build very large fabrics with maximum reliability. But just as critically, the fabric services are not independent of the network. You don't have a Domain Name Server (DNS) sitting off to the side of the environment, but rather these services are integrated and distributed within the fabric. This provides an incredible level of resiliency as every switch has a copy of the Name Server registry. Therefore, the loss of a single switch in the fabric doesn't cause the network to reconverge (where OSPF might take seconds to achieve a convergent view of the network to avoid potential loop creation for example) because all of the switches know the environment. If a known good alternate path exists from Initiator to Target, it will immediately be used and all non-affected traffic will never know the difference. This distributed environment also helps reduce the management time. As an example, a new switch added to the configuration will immediately learn its Name Server and Zoning databases from the existing fabric.

Since 1994 when Fibre Channel received its ANSI standard approval, companies have relied on and grown with it. As companies' storage requirements increased for higher availability, security and reliability, so did Fibre Channel fabric services to keep pace. Fibre Channel is a standards-driven protocol that addresses these concerns and today there are more than 100+ Industry Standards posted on Broadcom.com.

End-to-End Communication

The amount of data moving through the data center continues to grow faster than ever before, which makes it challenging to manage your infrastructure properly and extract business-level insights. On top of that, the company's success is now dependent on its ability to extract insights quickly. These demands drive up complexity. In fact, according to a January 2021 research report titled "Technology Spending Intention Survey," from IT and analyst firm ESG, 75% of organizations surveyed view their IT environments as more complex than they were two years ago. Adding to the complexity is how application owners access their data. Are the application owners wanting to get access through containers, virtual machines or a cloud? How does the storage keep track when applications are spun up and down so quickly?

The answer goes back to Fibre Channel being a collaborative protocol and evolving with the new requirements. Meaning the applications, server, storage and network vendors work together to address problems at the standards level to guarantee performance and reliability when application owners change the way they access storage. This is no small advantage. The ecosystem of Fibre Channel constantly tests hardware and standards-based functionality to ensure that the customer experience will not be the "we'll debug the new equipment in production," that IP storage environments can suffer from.

In most cases, the majority of storage networking issues can be addressed if the devices are aware of what is happening outside of their device. In other words, if the host, network and storage can talk and share events, then they each can take action to address issues that come up like speed mismatch, a failed cable or misconfigured MPIO path.

This is where the storage network plays the lead role in communication. The device in the middle sees almost everything and needs a way to share. For example, it's a standard practice to try and maximize your server resources by adding more and more virtual machines (VMs) because you have on average more compute power and storage capacity. But let's say one of those VMs starts to be overutilized and requests more storage resources than are available? The server side seems fine because the management tools say compute is fine so it must be a storage issue. The storage management tool says it's busy because of the host. So the finger-pointing starts.

Now with end-to-end communication between devices, the SAN switches can say tell the host that VM #xyz is the issue because it was not waiting for the storage response and was just dumping data. This communication is done through a fabric notifications mechanism that provides end devices with more information about events in the fabric. This includes notifications regarding link integrity issues, delivery notification issues and congestion issues. This means that instead of the management software trying to interpret what happened on the edge, the entire ecosystem (server, switch, storage) can be engaged in both identifying and correcting the problem.

What does this mean in context to Fibre Channel? This means that the SAN switch can send a notification to the HBA to address too many writes coming from the host. The HBA can take action by slowing down the traffic coming from that VM by throttling its performance.

This type of communication across Fibre Channel devices is only one aspect of reliability and leads me to the last topic.

Reliability

The reality of IT is that the lifecycle of infrastructure is based on getting the most use out of every device. This means that servers and storage purchased 3-5 years ago are still in use while new servers and storage are still being added. By connecting all of the elements to a network the new equipment may be dropped in next to the existing equipment.

This sounds like a great idea on paper but the intermix of multiple generations of SAN technology can frequently cause network issues. This is true regardless of the type of network as virtually no customer has an environment (Ethernet or Fibre Channel) where all of the servers are the same generation with the same versions of network interface controllers (NICs) and HBAs or where the storage elements are all the same model and generation. In fact, in networks that don't use buffer credits, the issue is even more significant as the mismatch in performance between generations will almost invariably lead to levels of congestion and the recovery mechanisms in networks where buffer-to-buffer credits are not in use take more time.

A simple way to think of this is that in a buffer-to-buffer credit environment data is never forwarded unless there is space to receive it. Whereas a TCP network will put the data on the wire and it is incumbent upon an upper-layer protocol at the end target to understand that data was lost and send a message back for recovery (a much longer relative wait time). Additionally, on packet loss, the congestion windowing algorithm in TCP will generally cut the throughput by 50% as a start point to begin bringing the traffic back under control which impacts the performance.

A standard definition of congestion, regardless of the protocol, would be when the rate of frames entering the fabric exceeds the rate of frames exiting the fabric. A simple solution to address this would be to tell applications to stop! That's not very customer-centric though. The goal of infrastructure should be that the application user never even notices when the infrastructure hits an issue and has to work around it. So, in the true spirit of collaboration across the Fibre Channel industry, they came up with a solution. The industry worked together to address reliability issues include features like:

- **Buffer credit** – Prevents a device from overrunning its peer
- **Flow control** – Paces the rate that devices can send data
- **Error detection/resource allocation** – Provides a mechanism for handling failing or misbehaving devices

Many network types indeed have recovery mechanisms. TCP/IP for instance will notice that packets were dropped and retransmit them. Fibre Channel by comparison uses the buffer-to-buffer credit mechanism to know in advance that there is space for the data about to be sent. Effectively, don't drop the traffic in the first place (it could be noted that the Peripheral Component Interconnect and InfiniBand protocols are also "buffer-to-buffer" credit mechanisms). So when typical Ethernet TCP/IP environments engage congestion management and congestion windowing reduces the traffic by 50%, the Fibre Channel environment will simply continue to process.

Summary

Companies, and specifically IT organizations, are facing numerous challenges when it comes to managing increasingly complex storage infrastructures. The storage network should not be adding to the stress. While no network is perfect and no single tool is ever perfect for every conceivable environment, the Fibre Channel storage area network continues to be the global workhorse when the requirements are high performance, lossless, low latency, reliable, time deterministic and secure delivery of storage traffic in the data center. To be fair, the same can be said about TCP/IP doing a brilliant job at delivering data in an unreliable physical network.

I expect both networks to continue to hold their place of prominence when traffic is mission-critical. TCP/IP will always be the best choice for connecting the Internet of Things and Fibre Channel for storage connectivity.