

Fibre Channel Never Dies

INSIDE
SANS

Learn the differentiating features
of Fibre Channel

Explore how Fibre Channel compares
to other storage networking protocols

Understand why Fibre Channel
is here to stay

By Juan Tarrío

About Brocade

Brocade, a Broadcom Inc. Company, is the proven leader in Fibre Channel storage networks that serve as the foundation for virtualized, all-flash data centers. Brocade Fibre Channel solutions deliver innovative, high-performance networks that are highly resilient and easier to deploy, manage and scale for the most demanding environments. The network matters for storage, and Brocade Fibre Channel storage networking solutions are the most trusted, widely deployed network infrastructure for enterprise storage.

www.broadcom.com




Table of Contents

| | | |
|------------------|--|-----------|
| Chapter 1 | The Multiple Deaths of Fibre Channel..... | 2 |
| | A Little Backstory | 3 |
| | The Convergence Wars | 4 |
| | The Converged Promised Land | 5 |
| | The Cycle of Life (and Technology) | 5 |
| Chapter 2 | The Flash Revolution and the NVMe Era - Fibre Channel Plays a Role..... | 8 |
| | The Flash Revolution | 8 |
| | The SAN Déjà Vu..... | 11 |
| | Between a RoCE and a Hard Place | 13 |
| | Fibre Channel is Ready | 14 |
| Chapter 3 | The Fibre Channel Complexity Myth | 15 |
| | A Convenient Half Truth | 15 |
| | Brocade Simplified Deployment and Operation of Ethernet Networks | 16 |
| | Avoiding Loops..... | 16 |
| | Layer 3 Complicates Things Even Further | 17 |
| | Lossless Ethernet and Data Center Bridging | 18 |
| | Explicit Congestion Notification | 19 |
| | Will TCP Work? | 19 |
| | Provisioning a Fibre Channel Network | 21 |
| Chapter 4 | Purpose-Built for Storage? | 23 |
| | Fibre Channel: A Transport Protocol | 23 |
| | Reliable, Pro-Active Flow Control | 25 |
| | Virtual Channels | 25 |
| | Beyond Five-Nines Availability | 27 |
| | The Many Use Cases of Ethernet..... | 29 |
| Chapter 5 | Fibre Channel Into the Future | 31 |
| | The Rise of NVME-oF | 31 |
| | Next Generation Non-Volatile Memory | 32 |
| | Fibre Channel as Dominant NVMe Transport | 33 |
| | Who Will Take the Ethernet-Based Alternative Crown? | 34 |
| | Customers Prefer Fibre Channel | 35 |



Introduction

Fibre Channel was developed in the 1990s in order to externalize, consolidate, centralize and share storage resources between multiple servers, pioneering the notion of a Storage Area Network (SAN), and enabling large-scale virtualization and cloud. Since then, Fibre Channel has had multiple contenders over the following years and decades. From alternative networking protocols to alternative architectures altogether, many technologies have tried to replace Fibre Channel as the gold standard for storage connectivity to applications and services in the enterprise data center, but have failed.

In this book, I will guide you through a brief history of Fibre Channel's purported replacements and will analyze its current contenders in the face of new and unprecedented innovations in storage technology.

About This Book

This book will help you dispel some of the myths and half-truths about Fibre Channel and its alternative protocols and technologies, helping you better understand some of the key characteristics that make it the gold standard for storage connectivity to mission-critical applications in the most demanding data centers in the world. It will also help you understand why Fibre Channel is best positioned to continue to be the storage networking technology of choice as we transition to a world of next-generation all-flash storage with unprecedented performance capabilities.

Chapter 1

The Multiple Deaths of Fibre Channel

In This Chapter

- Fibre Channel resiliency to alternatives
 - iSCSI emerges as an alternative to Fibre Channel
 - The convergence wars: Fibre Channel over Ethernet (FCoE)
 - Hyperconverged Infrastructures and Software-Defined Storage
-

With the advent of HCI, SDS, NVMe, NVMe-oF and SCM, you may have wondered whether Fibre Channel has any future. If you are reading this book, I'm sure you've read in the past that Fibre Channel is dead. Probably more than once. So, is Fibre Channel dead? There's a short answer, and that answer is: **Absolutely not.**

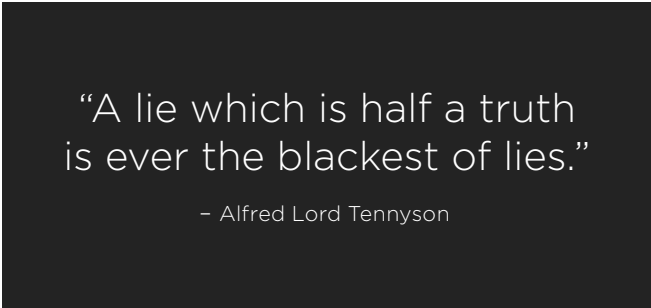
Let's discuss the reasons behind it. Fibre Channel is a resilient technology. No, I don't mean in the sense that you're thinking. It's resilient in that sense too, but what I'm driving at is its resistance to attacks from alternative technologies that position themselves as 'Fibre Channel Killers'. Truth is, even though Fibre Channel isn't that old, it's been declared 'dead' by proponents of such technologies several times since its birth in the '90s. It happened first with iSCSI technology circa 2002. Ironically, iSCSI has the honor of being one of the technologies that have purportedly killed Fibre Channel more than once.

A Little Backstory

Back in 2002, iSCSI worked at a whopping 1 Gbps, right on the heels of the industry's transition from fast Ethernet (which operated at 100 Mbps) to gigabit ethernet. At this time, Fibre Channel was just transitioning from 1 Gbps to 2 Gbps speeds—we were just releasing the Brocade SilkWorm 3800 and 3200 switches based on our 'Bloom' ASIC, and our first Fibre Channel director, the Brocade SilkWorm 12000. Why, then, would anyone consider iSCSI as a viable alternative to Fibre Channel for a Storage Area Network? The message is typically the same:

Fibre Channel is expensive, Fibre Channel is complex, it requires dedicated infrastructure and specialized skills; whereas IP/Ethernet is affordable, ubiquitous and everyone knows how to manage an Ethernet network, not to mention that everyone already has an Ethernet network!

These statements are all merely half-truths though. A wise man once said, "Tis better to have loved and lost than never to have loved at all." That same man also said, "A lie which is half a truth is ever the blackest of lies."



"A lie which is half a truth
is ever the blackest of lies."

– Alfred Lord Tennyson

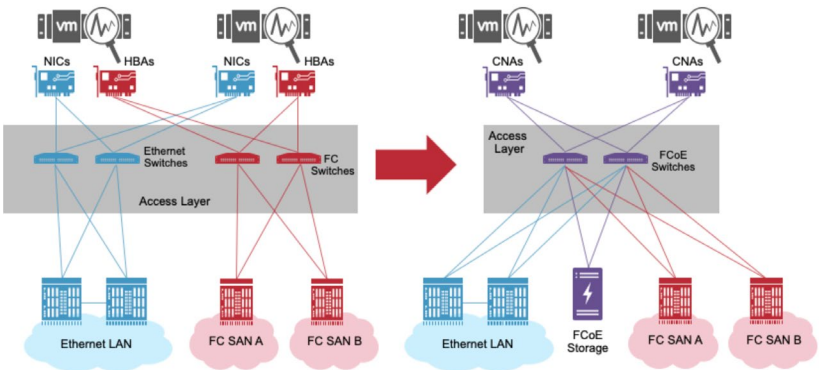
Of course, we all know what happened. Fibre Channel continued to be the dominant storage networking technology—particularly for mission-critical applications requiring the highest levels of performance and availability. iSCSI, on the other hand, became an alternative for some environments where performance and reliability weren't top requirements, but rather ubiquitous connectivity to a shared storage pool. After all, server virtualization was pushing hard to break out of the test and development environment, requiring all the physical servers to possess the ability to access a shared storage pool. And that's exactly what storage area networks had been born to facilitate.

The Convergence Wars

Enter 2008, and it was Fibre Channel over Ethernet's (FCoE) turn to claim that the death of Fibre Channel was imminent, at a time when the industry was riding the transition to 10 gigabit Ethernet. At this time, Fibre Channel 'only' delivered 8 gigabits of bandwidth. Once again, the proponents of convergence—the operative 'buzz word' at the time—would make the same arguments: Fibre Channel is expensive, Fibre Channel is complex, it requires dedicated infrastructure and specialized skills; whereas IP/Ethernet is affordable, ubiquitous and everyone knows how to manage an Ethernet network, not to mention that everyone already has an Ethernet network!

The sense of déjà vu was hard to shake off.

The Converged Promised Land



It was at this time when iSCSI renewed its assault on Fibre Channel, aided by new innovations in Ethernet technology, such as 10 Gbps speeds and Data Center Bridging (DCB). But once again, twelve years later, we all know what happened:

Fibre Channel continued to be the dominant storage networking technology—particularly for mission-critical applications that required the highest levels of performance and availability.

FCoE found its place where it made the most sense: at the edge of the network, where the consolidation of disparate I/O interfaces for storage and networking was highly desirable, while iSCSI continued to serve the environments where cost-effective connectivity was more important than performance and reliability.

To this day, many blade server architectures still use FCoE to save space and energy consumption within their chassis, but no one, not even Cisco, still believes that SANs and LANs will converge into One Big Network™, making everything Ethernet forever.

The Cycle of Life (and Technology)

If history does one thing, it repeats itself. All the time. As we ride the flash transition towards a world where NVMe will have replaced SCSI as the interface to storage, the proponents of Fibre Channel slayers are once again positioning several Ethernet-based technologies as alternatives and making the same arguments all over again: Fibre Channel is expensive, Fibre Channel is complex, it requires dedicated infrastructure and specialized skills; whereas IP/Ethernet is affordable, ubiquitous and everyone knows how to manage an Ethernet network, not to mention that everyone already has an Ethernet network!

In addition, over the last several years, we have seen Fibre Channel under attack not only by alternative protocols for traditional block-based SANs, but also by completely new and innovative shared storage architectures, claiming to entirely forgo the need to deploy a 'legacy' block-based storage network and purchase expensive, 'monolithic' storage arrays. Isn't that something? Of course, I'm talking about alternative architectures like Hyperconverged Infrastructures (HCI) and Software-Defined Storage (SDS).

HCI and SDS promise to deliver us from the chains of managing storage arrays or a storage network altogether. They are complementary technologies. In general, HCI just adds the hypervisor layer to an underlying SDS infrastructure, combining both compute and storage in building blocks or 'nodes' that can be stitched together over a network. The internal storage of all the different nodes is virtualized and made available to the hypervisors in a fully integrated way. In fact, these technologies are often called Virtual SANs (VSAN) or server SANs. Their main benefit is that they promise to integrate the provisioning of storage resources into the workflow of provisioning virtual machines, therefore greatly simplifying storage provisioning. They also promise to simplify operations altogether since there's no external storage array or a storage network to actually configure, provision, monitor and manage. On paper, these definitely sound like great benefits and are really appealing to IT admins and, more importantly, IT directors and CIOs.



There is a lot to be said about the advantages and disadvantages of ‘traditional’, ‘monolithic’ storage arrays connected to any form of storage network versus distributed server-based virtual SANs. However, that’s not the topic of this book. Ultimately, new technologies rarely replace and eradicate previous ones completely. They simply come and take their rightful place in the market: solving real customer problems, not those made up by vendors to try to prop up the technology they are trying to promote. Before the emergence of iSCSI, FCoE, HCI, SDS, and an assortment of other Three-Letter Acronyms (TLAs), Fibre Channel was the only option to deploy a shared storage pool, which is the foundation for virtualization and cloud. It was the only game in town. In fact, it was the technology created (by Brocade and other vendors) to enable just that, which certainly fueled its skyrocketing growth in the 2000s and early 2010s.

Choice is a good thing, though. Today, customers have a wealth of options to look at when they build an IT environment. They can choose the best protocol and/or architecture depending on their needs at any given time, or for any particular application. However, choice also means uncertainty and doubt. For this reason, it’s more important now than ever to have clear, concise, and accurate information—as opposed to half-truths—that will guide you to make an informed choice on the storage infrastructure you will deploy and build your IT services on.

Chapter 2

The Flash Revolution and the NVMe Era: Fibre Channel Plays a Role

In This Chapter

- Flash storage takes over the world
 - SCSI bows down to NVMe
 - NVMe outside the server: NVMe-oF
 - Ethernet-based alternatives and Fibre Channel's role with NVMe
-

Let me address the latest attack Fibre Channel is fending off from a pure 'protocol wars' point of view. This comes on the heels of one of the most interesting and disruptive transitions we have seen in the storage industry in the past few decades: the emergence of flash storage, the transition away from SCSI to NVMe and the imminent irruption of a new generation of non-volatile memory technology for storage that has come to be known as Storage-Class Memory (SCM).

The Flash Revolution

In today's world, we all know what flash storage is. We use it daily in our smartphones, tablets, and smartwatches, and we have it on our laptops. Who remembers the times when your laptop's hard disk drive would just keep spinning and making noise while your computer's performance ground to a halt? However, it wasn't that long ago when even these portable devices—or at least the ones that existed at the time—were still using traditional, magnetic, spinning HDDs—original iPod anyone?



In the data center, flash emerged at least a decade ago, but as new technologies that provide several orders of magnitude better performance than their predecessor often do, it came at a hefty price premium that relegated it to niche applications that were the only ones that could justify the extra cost, or to a sort of ‘cache’ layer to accelerate access to ‘hot’ data, while the bulk of the storage capacity continued to reside in spinning drives. These were called ‘hybrid’ storage arrays.


But as is usually the case with these sorts of technology transitions, as time went by and adoption ramped up, the prices came down, and it fueled a feedback loop of increased adoption and commoditized pricing, which paved the way for what we have today. We transitioned from installing solid-state drives (SSDs) that looked and behaved exactly like HDDs in traditional storage arrays to developing entirely new storage array architectures designed from the ground up to take advantage of the performance characteristics of flash storage—the so-called ‘All-Flash Arrays’ (AFAs), which left spinning drives behind for good.

Redesigning the architectures of storage arrays alone, however, has not been enough to take full advantage of all the performance that flash storage can provide. As we moved along this timeline, we realized that the traditional storage protocol we had been using to address these storage devices—our good and trustworthy SCSI—was ill-equipped to unleash the performance benefits of the transition to an all-flash storage world. Also, a new development in flash memory technology was on the horizon, one that would make it even more obvious that SCSI would simply not manage going forward.

The 'Small Computer System Interface' was developed in the late '70s and standardized in 1986. It was a parallel interface designed to be able to internally or externally connect a small number of diverse peripheral devices to a computer over a 'ribbon' cable. SCSI was not designed to connect just HDDs, but all sorts of computer peripherals such as floppy disks, scanners, printers, or CD drives. It was not designed with flash storage in mind, a technology that would not come into the world until many years later, so it was no surprise to anyone when we realized that it just couldn't keep up with the performance of the new storage devices it was being used to address.

That is why the industry came together under the Non-Volatile Memory Express (NVMe) organization to develop a new interface and protocol to—as stated by themselves in the organization's [website](#)—"fully expose the benefits of non-volatile memory in all types of computing environments from mobile to data center."

NVMe is not exactly new. Work on a specialized interface for accessing non-volatile memory (flash) began in late 2007 at the Intel Developer Forum, and from there the first standard was released in 2011, with commercial devices shipping in 2013. Since then, NVMe has consolidated itself as the new standard interface inside devices that use flash storage, including laptops and more recently desktop PCs and servers, typically running over a PCIe interface.



The SAN Déjà Vu

| | |
|--------------------|--|
| Ethernet Cards | Apple SSD Controller: |
| Fibre Channel | |
| FireWire | |
| Graphics/Displays | APPLE SSD SM0512L: |
| Memory | Capacity: 500,28 GB (500.277.788.672 bytes) |
| NVMeExpress | TRIM Support: Yes |
| PCI | Model: APPLE SSD SM0512L |
| Parallel SCSI | Revision: CXS5EA0Q |
| Power | Serial Number: C02818200DSHRFM1X |
| Printers | Link Width: x4 |
| SAS | Link Speed: 8.0 GT/s |
| SATA/SATA Express | Detachable Drive: No |
| SPI | BSD Name: disk0 |
| Storage | Partition Map Type: GPT (GUID Partition Table) |
| | Removable Media: No |
| | S.M.A.R.T. status: Verified |

An NVMe SSD inside an Apple MacBook Pro

But just as in the late '90s with SCSI and internal server storage, the industry started to realize that there could be serious benefits to moving the NVMe storage outside of the servers, centralizing it, consolidating it, and accessing it over a networked interface.

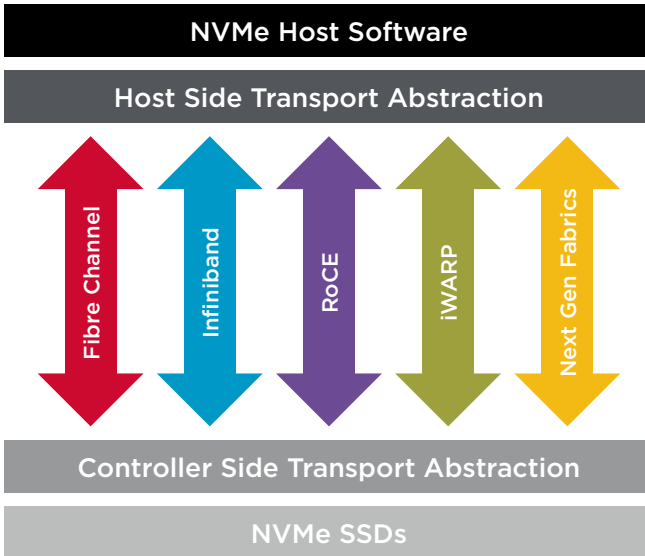
Once again, the sense of déjà vu is hard to shake off. This time around, the NVM Express organization itself had already anticipated this and published in 2016 a specification titled NVMe over Fabrics (NVMe-oF), in which they detailed how NVMe could be transported over “any suitable storage fabric technology.”

The NVMe-oF specification is purposefully agnostic about the underlying fabric, but does lay out key characteristics the “ideal underlying network or fabric technology” should have, because, as per the specification itself, “obviously, transporting NVMe commands across a network requires special considerations over and above those that are determined for local, in-storage memory.”

Some of those key characteristics that the ‘ideal’ underlying fabric should have include “a reliable, credit-based flow control and delivery mechanism” that can “guarantee delivery at the hardware level without the need to drop frames or packets due to congestion,” the fact that the fabric should “impose no more than 10µs of latency end-to-end, including the switches,” or that “the fabric should be able to scale to tens of thousands of devices or more.”

All while remaining agnostic, the NVMe-oF specification discusses two distinct types of fabric technologies that could transport NVMe over a network: those based on Remote Direct Memory Access (RDMA) on one side and those not based on RDMA on the other. Among the former group are Infiniband—the ‘native’ networked RDMA fabric technology—and its Ethernet-based alternatives, RDMA over Converged Ethernet in its second version (RoCE v2), and Internet Wide Area RDMA Protocol (iWARP).

RDMA is a protocol that allows a host to directly access shared memory space on another host, typically part of a supercomputing cluster. To briefly understand what RoCE and iWARP are, we could say that RoCE is to Infiniband what FCoE is to Fibre Channel—and when I say Fibre Channel in this context I mean SCSI over Fibre Channel—and iWARP is to Infiniband what iSCSI is to Fibre Channel, that is, two Ethernet-based alternatives to the dominant, native fabric in their space—which in the case of HPC is Infiniband—that have tried, to different degrees of success, to position themselves against a technology that is often described as complex, expensive and requiring specialized skills and dedicated infrastructure.



NVMe Over Fabrics Model

Between a RoCE and a Hard Place

Since the time leading to the release of the NVMe-oF specification in 2016, the proponents of Ethernet-based options to transport NVMe, particularly Mellanox (the main RDMA vendor proposing RoCE) went full-force on a marketing campaign to make it seem like RoCE was the only viable (or even existing) technology that could be used to connect NVMe devices over a fabric, and to claim that (once again) Fibre Channel was, yeah, you got it... dead.

From [blog posts in reputable publications](#) making the bold claim, to [blog posts in their own website](#) making the same claim, to [blog posts announcing the release of the NVMe-oF specification](#) that completely ignore Fibre Channel other than as a side mention as part of the author's experience to even [much more recent blog posts](#) trying to pretend that Fibre Channel doesn't even exist with statements such as "Simply, NVMe-oF stands for NVME over Fabrics [...] NVMe over Fabrics is essentially NVMe over RDMA".

In other forums, they would bring up arcane topics such as zero-copy that non-expert audiences knew nothing about and implied that it was something essential to the high performance of NVMe-oF—it is—and that only RDMA could provide—while ignoring the fact that Fibre Channel has supported it since the day the technology was invented.

It's no wonder, then, how RDMA—and RoCE, in particular—took the spotlight and most of the media and analyst pundit attention when it came to NVMe-oF and it seemed like, once again, Fibre Channel was an old, legacy technology that would not be there to support the new, exciting innovations in storage that were coming to market.

The message was the same it had always been: Fibre Channel is expensive, Fibre Channel is complex, it requires dedicated infrastructure and specialized skills; whereas IP/Ethernet is affordable, ubiquitous and everyone knows how to manage an Ethernet network, not to mention that everyone already has an Ethernet network!

Fibre Channel is Ready

In any case, Brocade wasn't dormant when it came to actual product research and development and the emergence of NVMe-oF in the technology landscape didn't go unnoticed. For one, little development was required to support running NVMe over Fibre Channel (NVMe/FC) on Brocade Gen 5 (16 Gbps) and Gen 6 (32 and 128 Gbps) switches and directors.

Technically, no development was required whatsoever just to be able to switch frames containing NVMe data, since Fibre Channel was developed as a transport protocol and NVMe is just another upper-level protocol (ULP) that is mapped onto it like SCSI or FICON—which is essentially ESCON running over Fibre Channel. An NVMe/FC frame is no different than a SCSI/FC or an 'ESCON/FC' frame from a Fibre Channel point of view, and technically any Brocade switch going back to the first generation would be able to switch it.

The only development required was for the name server to support devices registering NVMe as a supported ULP, which would enable NVMe initiator devices to easily discover available NVMe targets in the fabric, extending the benefits of the distributed fabric services to NVMe devices running over Fibre Channel and connecting to the same fabric as other devices running other ULPs.

This is one of the great advantages of running NVMe over a Fibre Channel fabric: it can easily coexist and be deployed alongside existing devices, whether open systems (SCSI) or mainframe (FICON), without requiring the deployment of new switches and without having to learn new ways of provisioning storage over an unknown fabric, therefore requiring the least amount of investment into either hardware or skills if you already run a Fibre Channel SAN.

Besides, running new NVMe devices alongside your existing install base of SCSI-based storage devices enables appealing use cases such as easy data migration from SCSI-based to newly-deployed NVMe-based arrays, snapshotting of existing databases running on SCSI-based arrays onto NVMe namespaces for fast big data analytics based on Machine Learning (ML) and Artificial Intelligence (AI), or extending SAN-based backup services to new NVMe-based storage devices.

Chapter 3

The Fibre Channel Complexity Myth

In This Chapter

- Is Ethernet/IP as simple as they tell you?
 - The challenges of Ethernet for storage
 - Fibre Channel advantages for storage provisioning
 - Fibre Channel is highly automated
-

A Convenient Half Truth

Let's address the claims that Fibre Channel is somehow an incredibly complex technology that requires terribly advanced skills to deploy and operate, while Ethernet is simple, affordable and everyone knows how to deploy and operate it simply by virtue of being Ethernet.

These are, as I mentioned earlier, half-truths. It's true that there is a much larger install base of Ethernet switch ports than there is of Fibre Channel switch ports out in the marketplace. Therefore, it's logical to assume that there are a lot more people trained and familiar with managing Ethernet networks than there are people capable of managing a Fibre Channel SAN. Even though there are more people trained in managing and configuring Ethernet networks, that shouldn't suggest that the technology is inherently simpler.

Brocade Simplified Deployment and Operation of Ethernet Networks

Before Brocade was acquired by Broadcom, we spent a number of years making a name for ourselves in the Ethernet/IP switching and routing industry. This followed our acquisition of Foundry Networks in 2009. The way we differentiated ourselves in the marketplace, particularly in the data center switching space, was by dramatically simplifying the deployment and operations of Ethernet networks. This was because—let’s face it—Ethernet networks have always been incredibly hard to manage and incredibly laborious to configure. In an Ethernet network, every single switch port needs to be told exactly what to do, whether it’s an access port—to connect end devices—or a trunk port—to connect other switches and form a network.

In the latter case, you must manually specify which VLANs are allowed to be carried over said trunk port. These aren’t the only port properties that need to be specified manually; if a port is to be part of a LAG—because the bandwidth of a single link has never been enough to carry all the traffic you need between two switches—it must also be specified manually. Otherwise you might accidentally create a loop and your entire network could melt down. This is because Ethernet is a layer 2 flood-and-learn protocol and if a loop exists you have very big problems.

Avoiding Loops

But how do you avoid loops in the first place? You would have to manually configure STP and give away half of your network’s bandwidth, and you must be willing to cope with seconds of network downtime if a link goes down. Or you could learn and manually configure its different variations like PVST or RSTP, but that will only make things slightly better. You could even try to avoid using STP altogether, but you’d have to use MLAG, which would require a manual configuration by you and would work completely differently for each vendor because a standard for MLAG doesn’t exist. Even then you’d better pray that it works well all the time, and make sure you have set up STP underneath it as a backup.



An apt visual representation of the Ethernet/IP protocol stack

Layer 3 Complicates Things Even Further

Once you're done setting your layer 2 domain, you have to start thinking about layer 3. Will you use IPv4 or IPv6? Do your switches have a big enough ARP cache? What routing protocol are you going to use? OSPF or BGP? Do you know how to properly configure ECMP for your routing protocol? How are you going to provide redundancy and high availability for your routing services? VRRP or HSRP? Do you need multicast services? Can you even spell IGMP? Do your switches support IGMP snooping? Do you know how to configure it? Do you want to completely get rid of L2 and STP and deploy a layer 3 fabric? Are you ready to manually assign an IP address to every single switch port? How are you going to make your virtualization layer believe it's still running on an L2 domain so that things like VM migrations work? Will you be using some kind of network virtualization technology based on VXLAN? Can your switches terminate VXLAN so you can extend your virtual L2 domains into the physical realm? Do you know how to configure that? Do you even know what VXLAN is?

And we haven't even addressed running storage traffic over the network!

Lossless Ethernet and Data Center Bridging

The proponents of SoE (Storage over Ethernet—yes, I made this one up) will just tell you, “Well, it’s just Ethernet! It’s simple! It’s interoperable! It’s scalable!”

Once again, we arrive at a number of half-truths. Paraphrasing the above quote from the NVMe-oF specification and adapting it for storage in general, we can safely say that “obviously, transporting storage across a network requires special considerations over and above those that are determined for local storage”, and therefore, ‘just Ethernet’ doesn’t cut it. In fact, if you run a flavor of SoE (like FCoE or NVMe/RoCE) that doesn’t rely on some upper-level protocol to ensure all packets get delivered to the destination (like TCP), then you’ll need to run a lossless Ethernet network.

For that you’ll need to run Data Center Bridging (DCB). Do all the switches in your network support DCB? Do your NIC cards on your servers support DCB? DCB relies on Priority-based Flow Control (PFC) to provide flow control at a granular level, so it’s only applied to the traffic that needs it (like storage) and not the traffic that would be hampered by it. However, it is still based on the use of pause frames and therefore it is a reactive mechanism that needs to wait for the receiver to detect that their buffer capacity is below a low threshold. Then, it will send a notification (PAUSE) to the transmitter for it to stop sending data and avoid dropping frames. This is unlike the proactive mechanism that technologies based on buffer-to-buffer (B2B) flow control deliver—yes, the ones that the “ideal” underlying network should support to run NVMe-oF.

Furthermore, your storage traffic is now sharing the entire network and its available bandwidth with all the rest of your network traffic. That means you need to configure Enhanced Transmission Selection (ETS), another part of DCB, to ensure that your storage traffic always has a minimum amount of guaranteed bandwidth available. However, this means that all your storage flows—potentially thousands of them—will be sharing a single ‘lane’, making it impossible to differentiate, protect, isolate or prioritize any of them.

Explicit Congestion Notification


In addition, RoCEv2 is routable, because it runs on top of UDP. While this means that it can potentially scale better than FCoE, because it can span across VLAN boundaries, it also means that the underlying flow control protocol that is supposed to guarantee frame delivery (PFC) is an L2 protocol itself and therefore cannot span across VLAN boundaries.

So, how do we ensure lossless delivery between end nodes in this case? By adding another piece to your “Jenga tower” and manually configuring Explicit Congestion Notification (ECN). Sounds simple, right? Well, first you need to make sure all your end-nodes and switches support it and that you actually know how to configure it or troubleshoot it if needed. Also, you need to make sure that it works reliably across devices from a variety of vendors. Now that I think about it, this doesn’t sound like “just Ethernet” to me.

Will TCP Work?

What if you run a flavor of SoE that relies on TCP for flow control and guaranteed delivery (mainly iSCSI, but also NVMe over iWARP or TCP)? Then you’ll have to deal with TCP’s well-known and widely accepted performance problems when packets are dropped and need to be resent (slow-start) as well as other issues. In fact, TCP is widely acknowledged to not be a good flow control protocol for low-latency, high-performance applications. That’s exactly what storage is.

Will TCP work well enough for several use cases? Of course, but that doesn’t mean it’s the right protocol for storage environments demanding reliability, deterministic low-latency and high performance. In fact, it has been so acknowledged by the storage industry that there are attempts to replace the TCP layer in storage for alternatives based on RDMA, such as the iSCSI Extensions for RDMA (iSER) or the SCSI RDMA Protocol (SRP). None of these have ever gained any significant traction, perhaps because of the added complexity of the RDMA layer, the need for specialized adapters called RDMA NICs (RNICs) or switches that support DCB and ECN, or perhaps even because they all fail to show significant performance benefits over traditional iSCSI over TCP—if they don’t perform even worse—as evidenced by the [RoCE Deployment Guide by Demartek](#).



Performance comparison between iSCSI and iSER

Whether you run your storage directly over Ethernet or over TCP/IP (or UDP/IP in the case of RoCEv2), there's still the issue of storage device discovery to be dealt with. An old comparison by EMC's Erik Smith on [his own personal blog](#) between FCoE, iSCSI, and Fibre Channel concluded that it takes a lot more configuration steps to provision storage on Ethernet-based fabrics. This is because there is no centralized name server or similar repository that can be used by end nodes for discovery, and therefore the storage resources need to be manually configured on every server somehow:

| | | FC vs. FCoE vs. iSCSI provisioning steps... | | | | | |
|---------|----------|---|---------------|----------|---------------|--------|----------|
| OS | Protocol | ...on host | ...on network | | ...on network | Total | |
| | | | no TDZ | with TDZ | | no TDZ | with TDZ |
| Windows | FC | 2 | 5 | 0 | 7 | 14 | 9 |
| | FCoE | 2 | 37 | 32 | 7 | 46 | 41 |
| | iSCSI | 19 | 7 | 7 | 7 | 33 | 33 |
| Linux | FC | 1 | 5 | 0 | 7 | 13 | 8 |
| | FCoE | 1 | 37 | 32 | 7 | 45 | 40 |
| | iSCSI | 14 | 7 | 7 | 7 | 28 | 28 |
| VMware | FC | 1 | 5 | 0 | 7 | 13 | 8 |
| | FCoE | 1 | 37 | 32 | 7 | 45 | 40 |
| | iSCSI | 23 | 7 | 7 | 7 | 37 | 37 |

Comparing the ease of storage resource provisioning between Fibre Channel, iSCSI and FCoE

I'll admit right from the start that I don't know if anything has been done in particular for NVMe/RoCE to aid in this—iSER is just iSCSI running on top of RoCE (or iWARP), but it's essentially just iSCSI—so perhaps things look a little better there? Maybe, but my suspicion is that they don't, and since RoCEv2 runs on IP, albeit with UDP instead of TCP, I'd wager you still have to manually enter the IP address of the target device in every initiator. This can be a huge operational burden in large environments with thousands of initiators.

While it is true that a service called iSNS (Internet Simple Name Server) for iSCSI exists—a service that can automate target device discovery for iSCSI initiators—the reality is that this is hardly ever implemented. Why? Because, to the best of my knowledge, there are no Ethernet switches that have an embedded iSNS server—Brocade had released embedded iSNS servers in our VDX switches just before the Broadcom acquisition—so users would have to deploy it in an external server. And there simply aren't any enterprise-class software iSNS implementations out there.

Suddenly, you have an overwhelming amount of different, interrelated protocols that create an incredibly complex protocol stack that you need to be able to provision, configure, manage, monitor and, in the event of something going wrong, troubleshoot. This isn't necessarily a bad thing, mind you. This is, in fact, part of the beauty of Ethernet/IP: it can serve a tremendous amount of purposes and support a tremendous amount of applications with varying degrees of service levels. However, to pretend even for a moment that Ethernet/IP somehow equates simplicity and ease of use or management is, as Alfred Lord Tennyson said, "ever the blackest of lies."

Provisioning a Fibre Channel Network

On the other hand, what do you need to do in order to provision a Fibre Channel network? In Fibre Channel, every switch port automatically detects what you connect to it and configures itself accordingly. This is true whether it's another switch or an end device. If it's another switch, it will automatically detect whether it's the first or subsequent link (Inter-Switch Link or ISL) between the same two switches. In the latter case, it will automatically figure out the best way to optimize load balancing between those ports: either at the physical layer with frame-based load balancing (only if you have Brocade switches) or at L2 with FSPF and exchange-based (I/O-level) load balancing. We could

describe Fibre Channel as a routed L2 network, and therefore there is no such thing as a loop, but rather multiple ways to get from one point to another.

If it is an end device, internal fabric services that run in a distributed fashion among all the switches in the fabric will help it determine which other devices it can communicate with or, in storage parlance, which target devices (storage) are available to each initiator (server). This is based on permissions configured centrally within the fabric by way of a technology called zoning. Not only that, the fabric will even enforce said permissions at a hardware level, automatically blocking and throwing away frames from unauthorized flows. The only prerequisite to be able to connect two Fibre Channel switches together is for them to have had a unique identifier (called Domain ID or DID) assigned to them by an administrator. But this operation is performed only once over a Fibre Channel switch's lifetime.

In general, the only thing that needs to be configured on an ongoing basis on a Fibre Channel SAN is zoning—which, by the way, is a process that can be automated with technologies like [target-driven zoning](#) or by using RESTful APIs. Why exactly this is considered 'complex' and requiring 'specialized skills' is anyone's guess. Of course, there are more advanced features that could be deployed on a Fibre Channel network, but few of these are really required for basic operations. There is a myriad of advanced monitoring, analytics, performance management, and proactive monitoring features that advanced users can take advantage of, specifically designed and developed for storage.

In fact, at Brocade we developed a technology that we came to call VCS (Virtual Cluster Switching). VCS borrowed a lot from the features that make Fibre Channel so easy to deploy and operate. Features such as auto-discovery of switches and fabric topology, or completely automated multi-pathing at the physical layer (which included our frame-based trunking that [made Ethernet networking experts say](#) "wow!"), as well as L2, by leveraging the same FSFP routing protocol running over a TRILL network. Our entire message was articulated around simplifying network deployment and operations, and it resonated well with customers.

Chapter 4

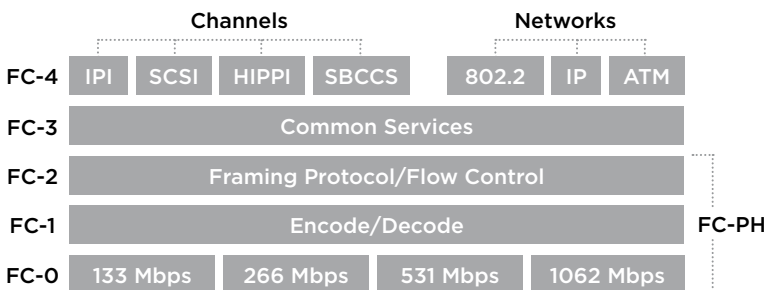
Purpose-built for Storage?

In This Chapter

- Fibre Channel as a transport protocol
 - Buffer-to-buffer flow control
 - Virtual Channels
 - Dual, air-gapped fabrics
 - Can Ethernet be good enough?
-

Fibre Channel: A Transport Protocol

One of the typical arguments you'll see Fibre Channel proponents make is that Fibre Channel is a technology that was purpose-built for storage. But, is that actually true? Certainly, companies like Brocade, Gadzoox, Ancor, Vixel and other startups did primarily focus on the storage use case for the products they were developing in the mid-'90s, but Fibre Channel was designed as a 'transport' protocol—although not in the way the OSI model defines the transport layer (L4)—that could transport other protocols on top of itself.



Fibre Channel protocol stack from the '90s

The highest layer of the Fibre Channel protocol stack, designated as FC-4, defines the mapping of different ULPs. As you can see from the diagram (taken from [this tutorial from 1994](#)) multiple ULPs were defined, including now obsolete ones like HIPPI or IPI, networking ones like ATM or IP, and some that are still in use today like SCSI or SBCCS (FICON).

Some of Brocade's earliest customers were media companies running video streaming applications on IP over Fibre Channel, which at the time outperformed Ethernet as it struggled to transition to Gigabit speeds with inefficient TCP/IP software stacks, while Fibre Channel had highly efficient hardware-based stacks and was transitioning from 1 Gbps to 2 Gbps. Brocade switches supported some specific features for IP over Fibre Channel and all Fibre Channel HBA vendors had IP drivers in addition to their SCSI drivers for all major operating systems.

However, it was ultimately the storage use case that propelled Fibre Channel to the position that it is in today, and other use cases slowly faded away. Now HBA vendors don't have IP drivers anymore and Brocade Fabric OS doesn't support the handful of IPoFC-specific features it once did. So even if Fibre Channel wasn't exclusively designed for storage, it might as well have, since it received decades and millions of dollars of R&D pretty much exclusively for the storage use case, both in open systems and mainframe environments.

That led to generation after generation of ASICs, operating system and management software versions that focused on more than just the speed bump that came with it. Dozens of features designed to more efficiently and reliably deliver thousands of concurrent mission-critical, high-performance, low-latency storage flows from their source to their destination while monitoring every single frame of every single flow in real-time, giving administrators the ability to measure I/O performance down to the individual storage LUN—or namespace ID (NSID) in the case of NVMe—even down to the VM level in virtualized environments, measuring not only throughput but also latencies, IOPS, first-response times, pending I/Os and many other storage-specific metrics that are incredibly valuable for the storage administrator.

Reliable, Proactive Flow Control

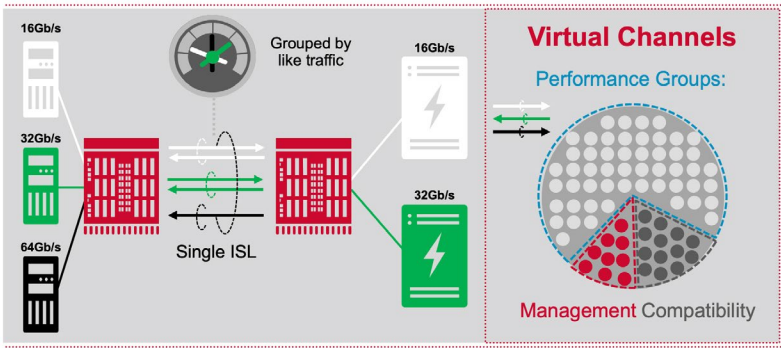
Fibre Channel implements a buffer-to-buffer flow control mechanism, by which every transmitter knows, upon link initialization, exactly how many buffers the receiver has available to hold frames before it processes them. The receiver grants the transmitter as many buffer 'credits' as buffers it has. The transmitter then keeps track of how many buffers the receiver has left by decrementing its credit counter by one every time it transmits a frame and incrementing it by one every time it receives a signal called Receiver Ready (R_RDY) that the receiver sends when it has processed a frame and freed a buffer.

This proactive mechanism ensures that the receiver is never overrun by an excess of frames from the transmitter that it cannot hold in its buffers and is therefore forced to discard. It is the mechanism the NVMe-oF specification deems 'ideal' for transporting NVMe traffic over a network, as it is the same mechanism that PCIe implements for internal NVMe storage inside a server. Buffer-to-buffer flow control has proven over decades to be a very good and reliable flow control mechanism for handling storage flows in a network.

Virtual Channels

Brocade developed a feature as part of our Fibre Channel ASICs that has been available since our first-generation ASIC called Virtual Channels (VCs). VCs automatically segment every ISL between two Brocade switches into several 'lanes', each with its dedicated buffer pool to provide independent flow control for each one of them. A small number of VCs are dedicated to the special traffic that exists to run the distributed fabric services (known as 'Class F' traffic) so that such traffic always has a dedicated maximum-priority lane even in cases of extreme congestion. That way, the fabric itself never becomes unstable because the switches can't communicate between each other.





Brocade Virtual Channels inside an ISL.

VCs enable the isolation, classification, protection, and prioritization of storage flows so that congestion events affecting one or some of them don't affect them all. If a storage device becomes unresponsive and turns into a slow-drain device its traffic flows start to experience increased latencies. If not acted upon, this will create 'back-pressure' on the network—something that is inherent to any flow-controlled network that cannot allow for frames to be dropped when there is congestion — and could potentially affect multiple unrelated 'victim' flows, impacting their application performance significantly and potentially causing serious consequences. Brocade Fibre Channel fabrics can automatically detect these increased latency conditions and 'quarantine' slow flows into low-priority VCs so that their performance degradation doesn't affect other storage flows.

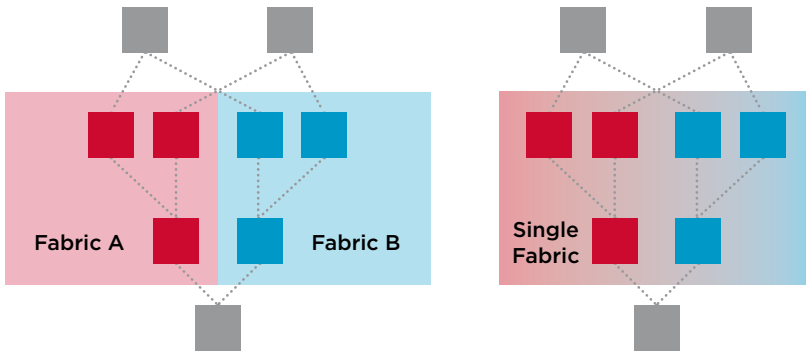
Virtual Channels technology has been available since our first-generation 'Stitch' ASIC that ran at 1 Gbps and continues to be available in our latest 'Condor 5' ASIC that supports Gen 7 Fibre Channel running at 64 Gbps. What we have done over the years is increase the number of VCs that are available for end-user traffic as well as develop software features that better take advantage of this technology, like Slow-Drain Device Quarantine (SDDQ), Quality of Service (QoS) or Traffic Optimizer (for Gen7).

Beyond Five-Nines Availability

As soon as we started taking storage outside of the servers and it was no longer connected to the CPU via internal buses, we realized as an industry how important it was to guarantee that an application or operating system's storage resource never became unavailable. I'm pretty sure you're aware of what happens when you disconnect a computer's hard drive (or SSD these days) while it's running. Then you can imagine the consequences when this happens to thousands of VMs that are running out of a SAN-attached array, or to a mission-critical application if its database becomes unavailable during operation. Not only will the application or the entire operating system of the server or VM crash, but there could easily be data corruption that would lead to extended periods of downtime with nefarious consequences, including going out of business.

For this reason, we started to develop technologies and best practices to ensure that there was never any single point of failure (SPoF) in a networked storage environment. In addition to redundant disks drives inside the storage arrays with data mirrored or striped across multiple drives to be able to withstand individual drive failure—a technology known as RAID—redundant controllers on the storage arrays paired with redundant adapters on the servers and a multipathing driver on the operating system, organizations started to deploy what has come to be known as 'dual fabrics', that is, two separate, completely air-gapped, no-single-cable-between-them storage fabrics so that any failure event on one of them could not, under any circumstance, ever, affect the other fabric.

This is only possible through complete physical isolation of the two fabrics, and that is why it is of utmost importance that these fabrics are physically air-gapped. Otherwise, no matter how much redundancy you build into it, you still have a single fabric, and therefore you have a potential single point of failure. Cynics would claim that this was all a ploy by greedy Fibre Channel vendors to convince customers that they needed to buy double the equipment. Of course, this argument can be easily refuted by showing that two different types of redundant networks can be built with the same number of devices, as in the diagram below.



Redundancy in Fibre Channel (left) and Ethernet/IP (right)

Does this mean that Ethernet/IP networks are not highly available? No. Does this mean that Ethernet/IP networks cannot be built with dual, air-gapped fabric redundancy, and therefore can never be as highly available as a dual, air-gapped Fibre Channel storage network? In a way it does. Ethernet/IP networks need to support a wide variety of use cases, the primary one being supporting TCP/IP communications between applications, clients and servers, and devices inside the data center and the outside world, or between a company's campus network and the internet.

Redundant links between Ethernet switches are based on LAG, (remember that there can be no loops) which in its inception, could only work between a single source switch and a single destination switch. Over the years, technologies like MLAG—not a single technology or standard, rather a myriad of vendor-specific proprietary implementations—were developed to overcome this limitation. MLAG technologies are based on making two switches behave as one, which requires them to be configured to do so by generally complex and laborious configuration steps. They often also require dedicating links between them for heartbeat and synchronization purposes.

Redundancy at the IP layer between an end device and the network requires two adapters on the device to 'team-up' (NIC teaming) and behave like one. They then present a single IP interface with a single IP address to the network. Technically, you could run redundant air-gapped networks in Ethernet/IP if you had a dedicated network just for storage, but that is hardly ever the case. Remember that one of the arguments for using Ethernet/IP for storage is precisely that you're not supposed to require dedicated infrastructure.

Similarly, specialized storage-focused performance monitoring, analytics, and troubleshooting tools don't exist for Ethernet-based storage networks not because Ethernet is inherently inferior to Fibre Channel and these kinds of tools cannot possibly exist, but because there hasn't been enough demand from customers or R&D time and money spent by vendors to develop said tools. Does this mean they can never exist? Of course not. But how likely is it for them to be developed, given that there is no single Ethernet switch vendor that is exclusively focused on storage?

The Many Use Cases of Ethernet

Once again, Ethernet (and Ethernet vendors) must support an incredibly wide variety of applications and use cases. Storage is only one of them, and not precisely one of the most important ones when it comes to port shipments and revenue, so it is unlikely that anyone will invest the time and money to develop them. Instead, customers will be left to use general-purpose performance monitoring, analytics, and troubleshooting tools that aren't designed for storage—and therefore don't provide the metrics that storage administrators require—leaving them unable to understand the behavior of storage flows or react fast enough when something is amiss.



Congestion in a network is not all that different from congestion in a road.

Similar things could be argued about how well Ethernet-based storage networks deal with the coexistence of thousands of storage flows, how they deal with congestion and backpressure, based on the flow control mechanism being used—whether it's PFC alone, PCF in combination with ECN, or TCP with or without ECN. In other words, how reliably they can deliver storage flows from source to destination.

Attempts at improving this are made now and then in the Ethernet space, like when DCB was developed to support FCoE, or the current DCTCP (Data Center TCP) initiative, which consists of new enhancements to ECN and TCP to improve exactly this in datacenter environments. Ethernet could even adopt buffer-to-buffer flow control, if needed—and if rumors are to be believed, this was proposed by a vendor when the industry was working on DCB and FCoE but was rejected. But the reality is that little R&D time and money is spent on storage use cases for Ethernet, as they remain a drop in the bucket of the Ethernet market.

When it comes to performance, we have recently proven with a [third-party validated report](#) that Ethernet/IP-based storage technologies—iSCSI in particular—simply can't take full advantage of the performance of modern all-flash storage arrays, or even fully utilize the network technology's nominal link bandwidth. If you are going to spend significant amounts of money on a high-performance all-flash array, you're going to want to take full advantage of your investment.

The question then becomes, how much high availability is enough? How much does a single percentage of application downtime cost? Not all applications are the same. They don't all require five-nines of availability, and there can be many use cases for which Ethernet/IP and the redundancy it can provide is good enough. Likewise, how much performance is enough? Not every application you run is going to require the highest levels of performance, or microsecond-level response times. There will be many applications where Ethernet/IP and the performance it can provide is sufficient. The same can be said about reliable delivery and congestion tolerance. In summary... how good is good enough?

There is no easy answer to this question, and—in the vast majority of cases—the answer will be “it depends”. Each organization is going to have to decide what qualifies as ‘good enough’ for their applications, and there won't be a one-size-fits-all answer anyway. Some of the services you will offer will require the highest levels of availability, some will require the best performance possible, and for many others it will be all about cost-effectiveness and not necessarily performance. What is important is that you understand what each technology can offer so you can make well-informed decisions and choose the right option for each of your applications or services—without half-truths. And it doesn't have to be a single solution for everything; it is perfectly possible for more than one storage infrastructure solution to coexist in your data center, each offering what they are best for.

Chapter 5

Fibre Channel Into the Future

In This Chapter

- NVMe over Fabrics takes off
 - Storage-Class Memory emerges
 - Who will rule among Ethernet-based NVMe-oF alternatives?
 - Customers prefer Fibre Channel
-

The Rise of NVME-oF

After having explored the set of capabilities that make Fibre Channel the technology best suited to support the highest demanding storage environments, it's time to look at the future and think about what can be expected of the technology going forward.

We are right at the moment when NVMe over Fabrics is going to take off in a significant way. Fibre Channel switch vendors—both Brocade and Cisco—already support NVMe/FC in our Gen 5 and Gen 6 switches and directors, as well as HBA vendors. Even all-flash array vendors have started to release arrays that support NVMe/FC on the front-end host ports, with several major vendors with solutions already in the market and many more to come in the coming months. This enables end-to-end NVMe/FC from the server, through the fabric, and to storage.

We have shown, with real-world test results and [in collaboration with Demartek, Emulex, and NetApp](#), that NVMe/FC can provide as much as a 58% performance improvement in terms of IOPS over SCSI/FC. We have also shown it can provide as much as a 34% reduction in latency, proving there are real benefits to be obtained from adopting this new technology.

Other tests performed [in collaboration with ESG, IBM and Emulex](#) have shown that NVMe/FC can deliver up to a 64% reduction in CPU utilization, which can bring significant savings, proving that the benefits of NVMe go beyond just performance.

Next Generation Non-Volatile Memory

This performance gap will only widen when current NAND-based flash technology gives way to next-generation non-volatile memory products like Intel Optane. Based on Intel's and Micron's co-developed 3D XPoint memory technology, products like Intel Optane are coming to be known as Storage-Class Memory or Persistent Memory, depending on whether the revolutionary memory technology is used as faster flash storage or as almost-as-fast-as-RAM non-volatile DIMMs (NVDIMMs) to complement and/or replace DRAM.

There is a lot that can be said about this new technology and the new use cases it will enable, including how applications could be rearchitected to take advantage of much faster storage access or non-volatile memory at nearly the speed of DRAM, but that is outside the scope of this book. Suffice it to say that when solely used as a replacement for existing NAND-based flash storage, it will provide a performance improvement of the same magnitude as the transition from spinning disks to current flash storage provided. This means it will place an even bigger performance and reliability burden on whatever network is used to transport it once it inevitably makes its way out of the server.

Fibre Channel as Dominant NVMe Transport

After an initial slow marketing start from Fibre Channel, the industry—not just experts, pundits, and vendors, but, more importantly, customers—have started to realize that Fibre Channel is best positioned to become the dominant transport for NVMe. The attention is starting to shift.

Not only is it best positioned from a performance, reliability, availability, and existence of a wealth of storage-specific tools point of view, it also provides the smoothest transition by being supported on the same infrastructure that is currently running most organizations' storage environments. This enables seamless deployments and migration without investing in new infrastructure or skills which, as I hope is clear now, aren't 'just Ethernet' skills.

Who Will Take the Ethernet-Based Alternative Crown?

I discussed NVMe over RoCE at length in one of the previous chapters, but as I also mentioned back then, the NVMe-oF specification outlined an additional Ethernet-based transport for NVMe: iWARP. As I explained, iWARP is another networked RDMA technology that is to Infiniband as an HPC protocol essentially what iSCSI is to SCSI over Fibre Channel as a storage protocol: just take the 'native' protocol (in this case Infiniband) and transport it over a TCP/IP network.

iWARP has limited market traction when it comes to alternatives to native Infiniband in HPC environments—where RDMA is actually necessary—and even less when it comes to the storage (NVMe) use case. No storage array vendor has ever expressed the intention of delivering support for NVMe over iWARP in their array host ports. For these reasons, iWARP is not expected to gain any momentum for NVMe.



It seemed pretty clear, given this picture, that RoCE v2 would come out on top as the winner among the Ethernet-based alternatives to Fibre Channel for NVMe-oF, and therefore as the official slayer of Fibre Channel (remember that this time it was “for real”).

However, things took a dramatic turn in the last couple of years and rumors of a new Ethernet-based NVMe transport started to emerge. Initially dubbed by some as ‘iNVMe’—because it is to NVMe/FC what iSCSI is to SCSI/FC—NVMe over TCP (NVMe/TCP) came into the scene backed by vendors such as Facebook and Intel, and later others like Dell EMC, NetApp or VMware (among others).

The idea behind NVMe/TCP is to transport the NVMe protocol directly over a TCP/IP network while completely doing away with the RDMA layer. Because, let’s face it, while flash storage is based on memory technology, we are still talking about storage, and the RDMA layer is completely unnecessary—remember that Fibre Channel has supported zero-copy from the start. In addition, this layer provides no value whatsoever and only adds complexity to the protocol stack.

Running NVMe directly over TCP means you can use any ‘mainstream’ Ethernet NIC without support for RDMA. And that’s precisely the point of NVMe/TCP: to be the equivalent to iSCSI in the SCSI/FC world, as a cost-effective and ubiquitous connectivity option for workloads that don’t have the performance and reliability requirements that demand Fibre Channel as the transport.

It is now generally believed in the industry that plain old TCP—with or without DCTCP—will be Fibre Channel’s biggest challenger for NVMe, while RoCE is widely perceived to be an arcane and complex technology for which it is even harder than for Fibre Channel to find people with the right skills to deploy.

This will leave us pretty much in the same situation we have been for many years where iSCSI was the main challenger for SCSI-based Fibre Channel networks, and we all know how that story went. There’s no reason to believe things will be any different between NVMe/FC and NVMe/TCP.

Customers Prefer Fibre Channel

Fibre Channel is therefore very well-positioned to take on this revolutionary transition that is starting to happen in the marketplace. If the flash transition has already acted as a boost to Fibre Channel port shipments and revenue over the past few years, we can only expect this trend to accelerate in the next few with NVMe/FC and SCM coming to the market. Most storage vendors are acknowledging this by bringing NVMe/FC to market on their storage arrays before any alternative based on Ethernet. Mainly because roughly 60-70% of all their all-flash arrays are already attached to a Fibre Channel SAN, and because it requires much less engineering effort to support NVMe/FC than any other alternative.

Plus, they're in a wait-and-see attitude with regards to which Ethernet-based NVMe alternative becomes favored by the market. Customers are realizing that they can deploy NVMe/FC on their existing SAN environments with near-zero risk, providing the most seamless transition to this new and exciting technology. Customers also realize they don't have to learn new ways to provision storage and can just leverage all the management, monitoring, analytics, and troubleshooting tools they know and love. Well, maybe love is a strong word.

I've said it before and I'll say it again, choice is good. When customers are presented with choice for their mission-critical storage environments, they keep choosing Fibre Channel.




So, is it true?

Is Fibre Channel dead once again as we find ourselves in one of the most exciting technological transitions for the storage market in the last few decades?

You better believe it isn't.
It's just getting started.

Key Takeaways

I hope that, after reading this book, you are better prepared to understand the unique features that Fibre Channel brings to the table and how they compare with what other storage networking protocols or alternative storage infrastructure technologies offer, so you are better prepared to make the right decision for your business.



Transform Your Network Into an Autonomous SAN with Brocade Gen 7

Brocade Gen 7 brings together hardware and software to create new dimensions of performance and reliability. Through self-learning, self-optimizing, and self-healing capabilities, Brocade Gen 7 delivers the intelligence your infrastructure needs to enable a faster, more efficient, and more resilient network.



Maximizes NVMe and scales devices with 2X performance and 50% lower latency



Transforms millions of data points into actionable intelligence



Resolves issues without intervention

Learn More
go.broadcom.com/fc-networking

About the Author

Juan Tarrío is a veteran Principal Architect for Data Center solutions at Broadcom with over 20 years of experience in Fibre Channel Storage Area Network (SAN) technologies. Juan has worked in a number of different areas since his start at Brocade in 2002, including professional services delivery, presales and product marketing. He was the lead technical evangelist for Brocade's Ethernet Fabric, IP Fabric, IP Storage and Software Defined Networking (SDN) technologies, and today he focuses on helping enterprises understand the different storage infrastructure technologies and how they meet their business needs.

Inside...

Learn the differentiating features of Fibre Channel

Explore how Fibre Channel compares to other storage networking protocols

Understand why Fibre Channel is here to stay