

## Apple and Cisco Get Down to Business

The latest development in easy collaboration and effortless business mobility.

I'd just returned after a week away from the office, but I hadn't yet turned on my Cisco Spark notifications.

While on my post-vacation grocery run, someone was urgently trying to reach me for a simple yes-or-no answer. However, when I received this Spark call on my iPhone, just as I would any other cell call, I was able to do far more than answer a basic question. I chatted with my colleague, viewed the document that was sent, and offered my suggestions.

The best part? I was able to accomplish all of that without any phone set-up whatsoever.

Today, 26 percent of in-office business calls are made on mobile devices. And that number is rapidly rising! Consequently, organizations are quickly embracing business mobility so they can empower their employees while increasing both efficiency and productivity. But it's not as simple as it may sound. Businesses must ensure that the experience is flawless for the end user, and easily manageable for the IT department.

Apple and Cisco, fortunately, are now making that happen.

The Apple-Cisco partnership is focusing on three things. One, optimizing Wi-Fi connectivity. Two, prioritizing business apps. And, three, integrating voice with collaboration.

Cisco Spark—the platform combining messages, meetings, and calls—now uses Apple's CallKit and SiriKit framework to integrate with iOS 10. So, not only can you answer Spark calls, but you can also place them using your Contacts, Recents, and Spark History lists, all from your lock screen. And that's just the beginning. Other call features like call waiting, call



prioritizing, Siri, and Bluetooth are easily accessible as well. Plus, whenever you want to collaborate in an even more integrated way, you can easily transfer calls to a Cisco Spark connected endpoint.

All of this, of course, is terrific news for mobile device users like you and me. Meanwhile, the IT teams who facilitate our collaborative interactions want these calls running through the corporate voice environment for better control, compliance, and cost reduction. Naturally, they want to maximize existing IT investments as well.

And they can.

The [Cisco Spark Hybrid Call Service](#) connects Cisco Spark to your Cisco Call Control. That means we don't have to use the cellular network when we travel and, as a result, can reduce our roaming charges. We can also make and receive business calls from anywhere simply by using our enterprise phone numbers. The partnership is also making sure business apps such as Cisco Spark maintain high-level priority on the optimized Wi-Fi network. The goal is to make our iOS 10 experiences feel intuitive and effortless. And it seems to be working. By the way, Fast Lane support is also available for WebEx and Jabber.

If you're like me, this is exactly what you've been hoping for. When I'm on my phone, I want to focus on getting work done, and not on setting everything up. Now I can do that.

If you're curious to know more, check out [Apple #TechTalks featuring Cisco](#) on iTunes. Or watch the Apple Cisco collaboration-centric TechWise video.

## Will New NVMe Storage Technologies Live Up to the Hype?

Ravi Kavuri  
VP, Data Management Group & Chief Architect  
NetApp

*This blog is part 1 of a four-part series that explains how Nonvolatile Memory Express (NVMe), NVMe over Fabrics (NVMe-oF), and new storage-class memory (SCM), also known as PMEM, are changing the game for data centers. For a deeper dive, download the white paper "New Frontiers in Solid State Storage."*

Any major new technology development is bound to generate buzz. Sometimes, it turns out to be a fad, and everyone wonders what all the fuss was about. Other times, you can look back and realize that it was just the beginning of a trend that would be with us for years. We're seeing this right now with a new wave of solid state technologies that are just starting to hit the market. And, while I hesitate to predict the future, it looks like they'll be firmly fixed in the long term trend camp. Why? Because I believe adopting these technologies the right way will redefine how we design and build a next generation data center infrastructure that delivers consistently low latency for a wide range of workloads, at the lowest cost, via a robust ecosystem of best of breed suppliers.

I'm talking, of course, about three new innovations:

- **Nonvolatile Memory Express (NVMe)**, a protocol that allows for fast access for direct- attached flash storage. NVMe is an evolutionary step towards exploiting inherent parallelism built into SSDs and other solid state technologies.
- **NVMe-oF**: NVMe-over-Fabrics (NVMe-oF) exposes the advantages of NVMe to a fabric connecting hosts with networked storage. With increased adoption of low latency, high bandwidth network fabric (Ethernet, IB, Fibre Channel), it is now possible to build infrastructure that extends the performance advantage of NVMe over standard fabrics to access low latency non-volatile persistent storage.
- **Storage-Class Memory (SCM)**, such as Intel 3D XPoint and Samsung Z-NAND media, connected over NVMe technology or even on the memory bus, enable device latencies about 10 times lower than today's NAND-based SSDs.



These technologies by themselves don't mean anything without the ecosystem software (operating system/hypervisor, drivers, and protocols) being optimized. The NVMe software ecosystem is being designed to take advantage of these low latency transport and media. The combination of ultra-fast media types & low latency access mechanisms enabled by RDMA capabilities enable new disaggregated data center system architectures that can be deployed on standard fabrics familiar to customers. There is a class of emerging applications that can take advantage of this ultra-low latency infrastructure, enabling customers to gain significant insights by shifting the focus from post process to real time handling of data.

### Where Will NVMe and SCM Be Implemented?

From a storage system perspective, NVMe-oF will be deployed in two contexts: front-end (from server to storage system) and back-end (from storage system to NVMe device). Together with the current FCP front-end and back-end SAS/SATA choices, we have a plethora of combinations that are possible. SCM media will initially be used as a read/write cache to provide significantly lower latency than today's NAND Flash SSDs. As the price of SCM media comes down and it becomes a viable option for more applications, you'll be able to create a pool of all SCM storage to deliver consistent low latency that is an order of magnitude faster than today's shared storage. NetApp will roll out some of these new technologies over time as they mature while protecting investment in existing storage technologies.

### What Results Can You Expect?

You can expect significant, tangible cost savings in your datacenter infrastructure, as these new NVMe, NVMe-oF and SCM technologies drive down CPU overhead in the I/O stack and allow external storage arrays to operate much more efficiently. This will allow for better utilization of both server compute and storage resources for a variety of workloads, including business intelligence, analytics, and data warehousing. Fewer CPU cores handling the same set of application workloads will lead to reduction in application licensing costs and higher application density.

Moreover, emerging applications with more stringent service level objectives will now be feasible in enterprise data centers. The next-generation of real-time analytics, machine learning, and artificial intelligence will demand real-time response from a highly parallel, ultra-low latency data management infrastructure.

### Looking Ahead

It's an exciting time in storage. At NetApp, we're playing a lead role in the development of these new technologies. Most importantly, we're making sure that our customers can capitalize on them without having to rip and replace their infrastructures, or sacrifice the NetApp data management features they rely on every day. In the future, you'll be able to use front-end NVMe-oF in NetApp's All Flash Array products to gain higher throughput at lower latencies, and use less CPU at the server and in the storage system. You'll also be able to use additional NVMe form factors and media types in our All Flash FAS (AFF) family to take advantage of the highest-performance SSD devices. Keep an eye on this space, as we explore the implications of these new innovations in upcoming blogs.

*Visit our booth at Flash Memory Summit:  
August 4 - 10, 2017 and attend our keynote: Creating  
the Fabric of a New Generation of Enterprise Apps.  
Thursday, August 10: 11:30AM - 12:00PM.*

## Three Simple Questions You Should Ask Every Flash Vendor re: NVMe

Ravi Kavuri  
VP, Data Management Group & Chief Architect  
NetApp

*This blog is part 4 of a four-part series that explains how Nonvolatile Memory Express (NVMe), NVMe over Fabrics (NVMe-oF), and new storage-class memory (SCM), also known as PMEM, are changing the game for data centers. For a deeper dive, download the white paper "[New Frontiers in Solid State Storage](#)."*

In the coming years, every major storage vendor will be rolling out its own solutions for NVMe, NVMe-oF, and SCM. But as with most things in life, the devil is in the details. Different vendors will implement new solid state technologies and architectures in very different ways.

As you evaluate your options, don't be afraid to ask tough questions. Make sure that in all the talk of amazing speeds and feeds, your vendor isn't pushing a solution that will create more problems than it solves.

### Put Flash Vendors to the Test

Most organizations won't overhaul their entire storage architectures overnight. They're more likely to bring in new technologies for specific, targeted storage. For example, organizations may initially deploy SCM only for the most performance sensitive applications, while continuing to use flash SSDs for the rest. The best way to do that is to be able to non-disruptively introduce these new solid state storage technologies as just another tier of storage that functions alongside your existing storage (both flash and HDDs).

Is that what your vendor is offering?  
Here are some key questions to find out:

#### 1. What will I have to give up to use NVMe and SCM capabilities?

You shouldn't have to sacrifice enterprise-class data management features and resiliency to use new storage technologies and create another silo.



You choose storage vendors not just for the latest technology, but because they offer mature and stable software, resiliency and redundancy, flexible data management and simple application integrations. If choosing a new storage technology means sacrificing any of that, tread carefully.

#### 2. How will this solution evolve?

Underlying hardware technologies can change quickly. Avoid solutions custom-built from the ground up for a specific new technology, or you'll have a hard time adding new capabilities over time. And run away from any vendor insisting that their custom-built solid-state media in a proprietary form factor can compete in even the medium-term with the inherent price, performance, and longevity of commodity-based solutions from giants like Intel, Samsung, and Toshiba. Don't get locked in. Instead, look for software-centric solutions that can continually incorporate new underlying hardware and storage media without making wholesale changes.

#### 3. How will this affect my existing environment?

You don't want to have to deploy an entirely separate silo for each new storage technology in your data center. Rather, you should be able to use mature clustering technology to not only introduce new technology non-disruptively but to provision and manage applications that adhere to your service levels regardless of the specific mix of underlying architectures and storage media.

Bottom line, as exciting as new storage innovations are, they don't change your core requirements for storage: high-performance, reliable access to your data, with simple management and operation. NVMe, NVMe-oF and SCM technologies really can bring amazing new capabilities. But only if they're designed to function as part of your real-world data center, not a science experiment.

### NetApp Aces This Test

At NetApp, we offer some of the fastest flash storage platforms in the industry, and we constantly release new features and capabilities for our products. But we do it with a focus on software innovation and scale-out architectures, so our customers can continually extend the value of their NetApp investments.

We're bringing this same approach to NVMe, NVMe-oF and SCM. As new technology innovations offer more performance, density and cost-efficiency, our customers will be able to take advantage of them easily, non-disruptively and as part of enterprise-grade storage platforms. Visit our booth at Flash Memory Summit: August 4 – 10, 2017 and attend our keynote: Creating the Fabric of a New Generation of Enterprise Apps. Thursday, August 10: 11:30AM – 12:00PM.

*Visit our booth at Flash Memory Summit:  
August 4 – 10, 2017 and attend our keynote:  
Creating the Fabric of a New Generation of Enterprise  
Apps. Thursday, August 10: 11:30AM – 12:00PM.*