AUGUST 2012 / VOL. 3 / NO. 4

NETWORK evolution

BUILDING THE INFRASTRUCTURE TO ENABLE THE CHANGING FACE OF IT

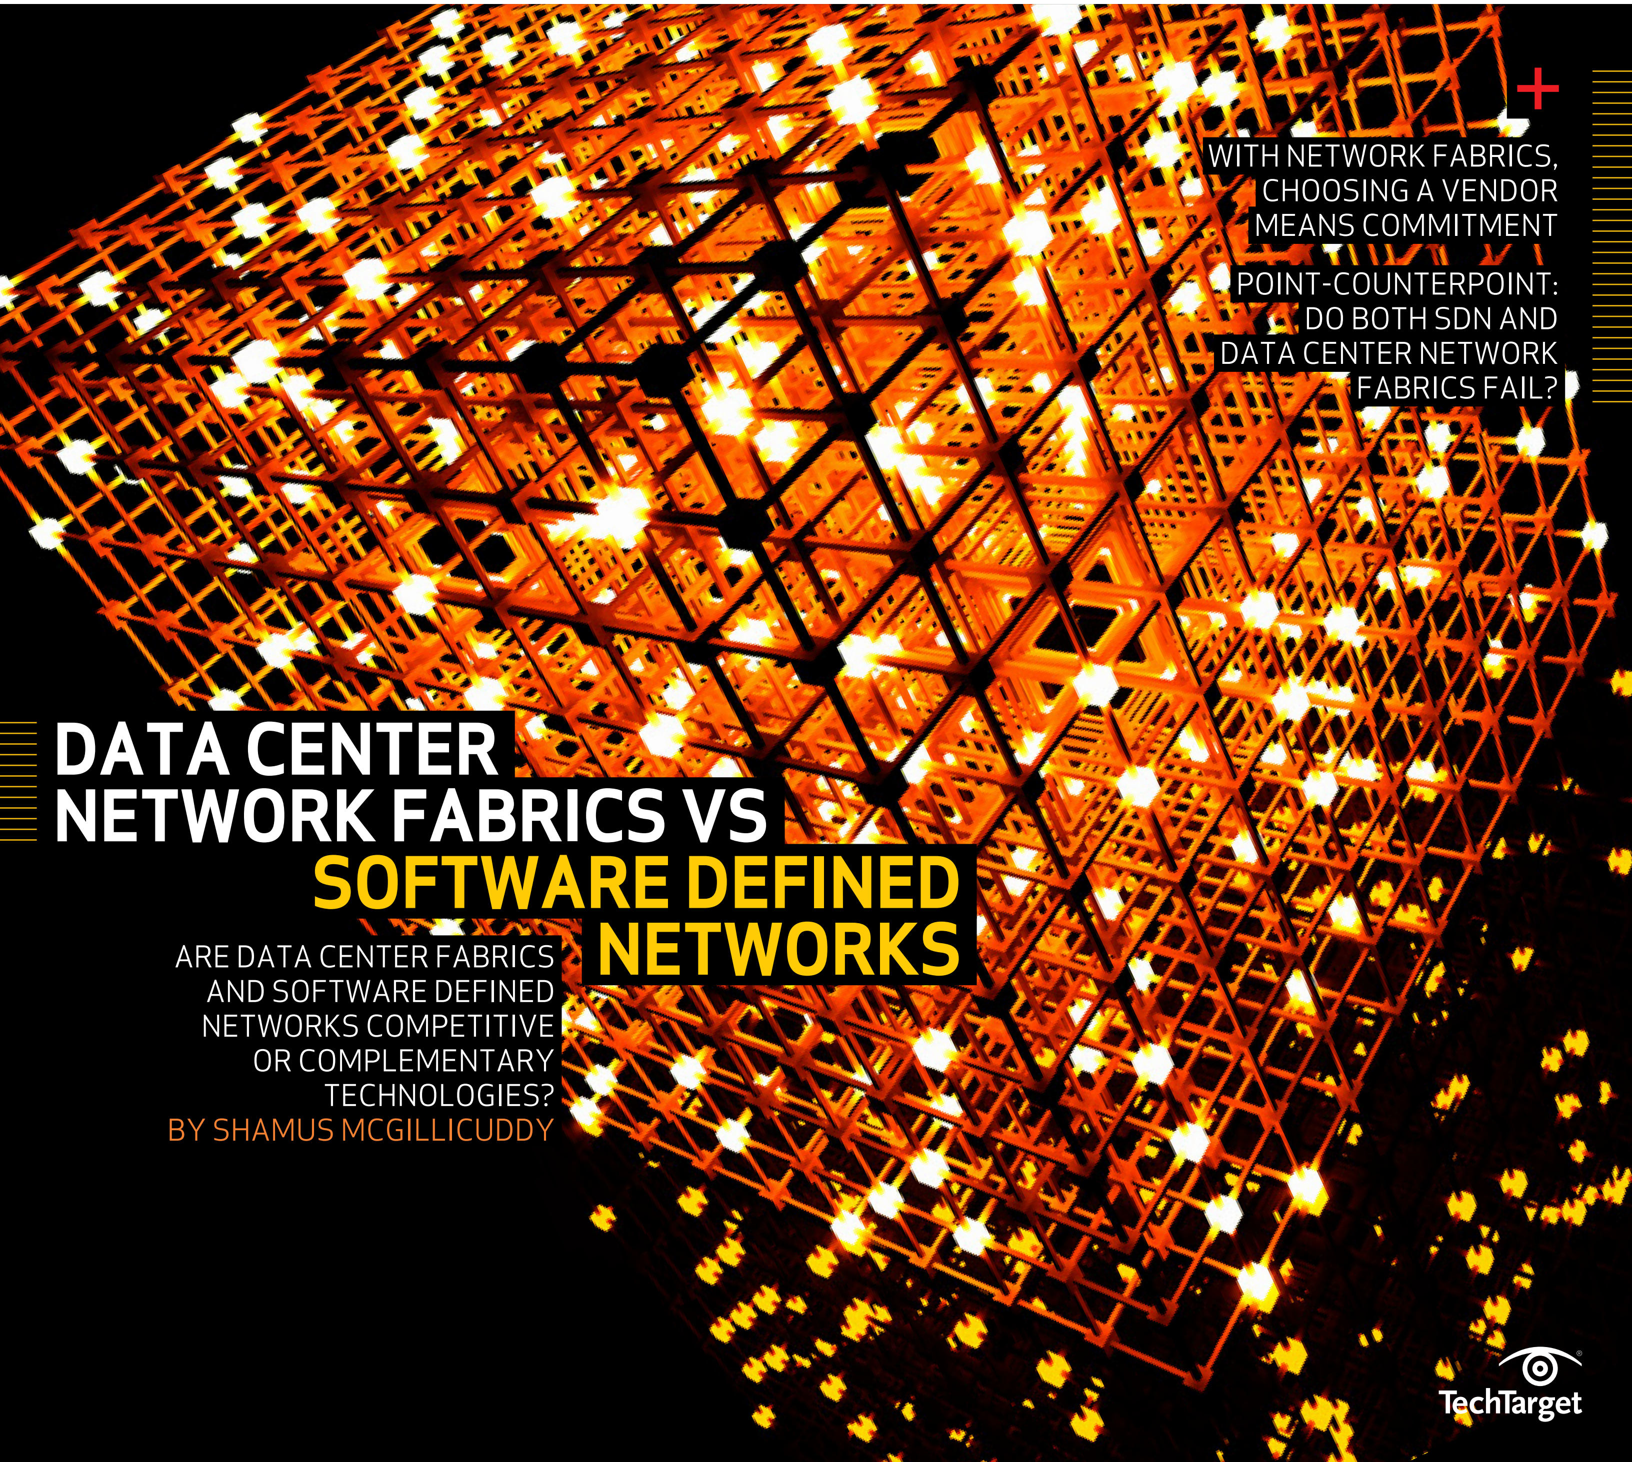

+ WITH NETWORK FABRICS, CHOOSING A VENDOR MEANS COMMITMENT

+ POINT-COUNTERPOINT: DO BOTH SDN AND DATA CENTER NETWORK FABRICS FAIL?

DATA CENTER NETWORK FABRICS VS SOFTWARE DEFINED NETWORKS

ARE DATA CENTER FABRICS AND SOFTWARE DEFINED NETWORKS COMPETITIVE OR COMPLEMENTARY TECHNOLOGIES?

BY SHAMUS MCGILLICUDDY

TechTarget

idealab

Where evolving network concepts come together

BY DAVID B. JACOBS

Deep Packet Inspection Tools: Proxy vs. Stream-Based

DEEP PACKET INSPECTION (DPI) tools have been mostly associated with service provider networks, but enterprise network managers are increasingly turning to the technology to better manage application performance and ensure a greater level of security.

DPI tools inspect the contents of a packet and determine performance based on which application layer protocol is in use. As such, DPI makes it possible to find, identify, classify, reroute or block packets with specific data or code payloads that conventional packet filtering cannot detect.

Packet inspection strategies can be broken into two categories: stream-based and proxy-based.

Stream-based inspection examines the data in each incoming packet as it arrives. If no threat is found, the packet is forwarded to its destination. Proxy-based inspection buffers the series of packets that make up a single transaction and inspects for threats after all packets have been received. Both stream- and proxy-based inspection techniques match data sequences against known threat signatures and also utilize heuristics to detect zero-day attacks.

> Packet inspection strategies can be broken into stream-based and proxy-based.

Critics of proxy-based DPI tools say that the volume of data pouring through protection devices (especially with increasing file sizes) makes it impossible for a proxy-based product to buffer all of the incoming traffic. What's more, they believe that buffering large files introduces unacceptable delays in application performance.

Meanwhile, critics of stream-based

technology say those tools aren't as thorough as proxy-based tools because it is impossible to detect threats without viewing the entire transaction. They add that stream-based products can decompress only basic compression techniques like .zip, while proxy-based products can decompress many techniques. Stream-based product vendors contend that their software can detect the characteristics of malware as they inspect packets one by one.

You can learn more about DPI vendors in this vendor comparison, where SearchNetworking runs down a sampling of vendors and their offerings. ■

BY LISA PHIFER

Using Wirless Network Bandwidth Monitoring To Stay Within Data Caps

GIVEN 3G AND 4G bandwidth caps and new cloud content synchronization services like iCloud, how can wireless broadband data users conduct network bandwidth monitoring and avoid exceeding limits?

Monitoring bandwidth usage has long been a challenge for wireless broadband users. Checking a carrier's website to eyeball usage-to-date is hardly a real-time solution, and new cloud apps are boosting usage to record highs. For some users, the first indication that something is amiss is an unexpectedly big bill. But figuring out which app or setting is the culprit can be tough.

Enterprises can deploy mobile expense management tools, using device-resident agents for monitoring bandwidth usage. Such tools can help enterprises enforce bandwidth caps by generating admin or user threshold notifications.

Individuals can accomplish part or all of this through native capabilities or third-party apps. Older Android, iPhone and iPad users can download free or paid network bandwidth monitoring apps from Apple's AppStore.

Android 4 (Ice Cream Sandwich) users can now just tap Settings/Data Usage to see this information graphed against any mobile data limit the user might want to set. This graph is followed by a per-application breakdown, making it easy to identify apps consuming the most mobile data. In fact, this can be a good way to spot cloud apps that are silently eating away at your bandwidth cap through automatic data synchronization.

Once you know that you're approaching your limit—and preferably which apps are busting your budget—other device settings can be used to

HOME

IDEA LAB

DATA CENTER NETWORK FABRICS VS. SDN

WITH NETWORK FABRICS, CHOOSING A VENDOR MEANS COMMITMENT

POINT-COUNTERPOINT

whittle away at overages. For example, on iOS 5 devices, administrators or users can turn off iCloud backup or document sync and photo streaming—or disable iCloud entirely. On Android devices, Settings /Accounts and Sync can disable auto-sync for apps that are rarely used or temporarily when roaming. Many individual apps also have settings to fine-tune network use by setting refresh intervals, disabling notifications, requiring manual sync when roaming (e.g., Exchange Active Sync v14), or warning when high-band-width apps are used over mobile broadband. ■

BY SHAMUS MCGILLICUDDY

VMware Networking CTO Describes SDN and Network Virtualization

ALTHOUGH VMWARE IS best known as a server virtualization company, its impact on the networking industry has been profound. VMware-based server consolidation has driven up bandwidth requirements; the embedded distributed virtual switch in VMware's hypvervisor software has virtualized the server access layer of most data centers; and the dynamic nature of virtual machines is driving much of the competitive innovation in the data-center networking market today.

During his keynote at Interop Las Vegas 2012, Steve Herrod, VMware CTO and senior vice president of R&D, introduced the concept of the software-defined data center, a term which is naturally evocative of software-defined networking (SDN). SearchNetworking spoke recently with **Allwyn Sequeira**, CTO and vice president of cloud, networking and security, to learn more about VMware's views on the software-defined data center, network virtualization and SDN.

VMware has started promoting the concept of the software-defined data center. Could you elaborate on this?

The thinking was, we had vSphere and life was good. Now we need to virtualize the rest of the data center. As we started to evolve our thinking toward networking, the software-defined-networking momentum also happened to be in play at the same time. Most of our concepts of network virtualization, and some of the notions of SDN, are very applicable across network, compute, storage and security. It's a natural evolution to take the SDN construct and parlay that into a larger construct, which is the software-defined data center. Over time, we expect to see less and less specialty hardware [in data centers] and more x86 hardware with specialty software.

The other theme is that the virtual machine [VM] as a container has stood up well and become the abstraction in the workload space, whether it be Amazon or VMware. The VM has become a unit of server virtualization and desktop virtualization. Likewise, what's left of the data center, the virtual data center, becomes a new container.

What do you mean when you speak of data centers being composed of x86 hardware with specialty software? Is this the VMware networking vision?

A broader statement is x86 and merchant silicon is clearly the trend that exists. My point there is that the Ciscos of the world have spent lot of time and energy on ASICs, but a lot of software and control is increasingly happening on the x86 portion. Most of [data center] gear is headed that way. If you look at an F5 [Networks] device, it's mostly x86s. Look at a Cisco ASA, where they have firewalls and load balancers. That's all x86-based. We have versions of Cisco routers working with embedded hypervisors running on x86 to enable a whole bunch of VMs to run on them. If you look at some Cisco devices, they run with a hypervisor embedded. Not many people recognize it, but an increasing part of hardware, even in the Cisco world, is x86-based and merchant silicon. So the better way to say it is that X86 and commercial off-the-shelf merchant silicon is going to be prevalent in data center servers and the network itself.

I think Cisco might disagree with you. It continues to argue that ASICs give it an edge in the data center.

If I were Cisco, that's what I would say. And if I were Cisco, I would still look at merchant silicon. If I purely based my franchise on merchant silicon, that really says the networking play moves to software. What I think they're saying is that all that merchant silicon and x86 is fine, but when you get to higher and higher consolidation, there is still a need to scale out with rows of x86-based, top-of-rack switches. You still need to come out to core. And at the heart of it, you still need very fast switching silicon there, the likes of which are only produced by the Ciscos and Junipers of the world. And that is a valid statement.

Any Cisco box has three components: input and output ports, software processing and then the fastpath. The fastpath—tagging and protocols—is where Cisco has always distinguished itself from its peers. Inside the very core of these boxes, having that fast and specialized silicon might make sense. But all the programmability and scale-out and manipulability will come out of the slowpath and control plane portions of these boxes. ■

DATA CENTER NETWORK FABRICS VS. SOFTWARE-DEFINED NETWORKS

Are data center fabrics and software defined networks competitive or complementary technologies?

BYSHAMUS MCGILLICUDDY

JUST LAST YEAR, data center network fabrics were the hot topic in networking. Flat, low-latency networks with any-to-any bandwidth promised to solve the networking problems of highly virtualized data centers. These fabrics would enable increased east-west traffic in—and free up bandwidth constraints created by—spanning tree protocol (STP), making networks responsive to server virtualization and resulting in easier-to-operate data center infrastructure.

Then something else happened—a tsunami of hype hit the networking industry in the form of OpenFlow and software defined networking (SDN). Startups emerged with claims that SDN could enable programmable networks in multivendor environments, solving many of the problems that expensive data center fabrics promised to fix, but for a whole lot less money.

Now network engineers and architects must sort through two separate hype machines. And they must ask whether data center fabrics and SDN are an either-or proposition, or two architectures that complement each other.

The answer to this question is hard to find, particularly because fabrics and SDNs are still very new to the market. Few enterprises have moved beyond the proof-of-concept stage with fabrics, and the availability of robust, commercial SDN products remains scarce. In fact, SDN is still evolving, and the use cases for it are still developing. What's more, every vendor has its own selfdefined data center fabric offering, but most are also cooking up SDN plans.

Proponents from either camp will tell you that their side can solve most data center networking needs. However, they'll also concede that there is room for both.

"The goals [of SDN and fabrics] sound the same, but abstracting the network to extract complexity is one thing, and then fundamentally simplifying the infrastructure is another thing. The two working together can be extremely powerful," said Mike Marcellin, vice president of product marketing and strategy at Juniper. "If you simply abstract the complexity like an SDN is trying to do, that can help. But if the fundamental architecture is still complex, if it's still brittle, if you have to keep throwing more devices at the problem to scale the network, then the root of the problem is still not addressed."

Dan Pitt, executive director of the Open Networking Foundation, which governs the development of OpenFlow, concedes Marcellin's point to a point.

"You could choose one or the other, or you could choose a combination," said Pitt. "But I think if a customer has a brand-new choice, a true SDN solution with its inherent benefits and with its many choices in how they procure and deploy it [that will be their choice]."

SDN IS NOT QUITE READY TO REPLACE INNOVATIVE HARDWARE

When SDN and OpenFlow hype first started taking off, some industry observers predicted that OpenFlow's ability to centralize the control plane of a network would commoditize network hardware. Switches and routers would become dumb boxes pushing packets back and forth, and the brains of a network would live on a server. However, reality has since started setting in.

"The less reasonable part of the [SDN] discussion is whether network companies are dead because young entrepreneurs can come along and buy a few parts from the silicon vendors and build stuff that is equivalent to what switch vendors do," said Peter Christy, principal analyst with Internet Research Group. "I think that part is unrealistic because high performance, highly reliable fabric that operates on 40 Gbps links is pretty interesting engineering."

Also, SDN remains more of a curiosity for enterprises, rather than a commercialized solution that can solve problems today. Google famously implemented OpenFlow on its WAN, thanks to a small army of internal engineers who have the expertise to build home-brewed Google gear. Similarly, SDN startups have helped some other Web-scale companies and cloud providers

HOME

IDEA LAB

DATA CENTER NETWORK FABRICS VS. SDN

WITH NETWORK FABRICS, CHOOSING A VENDOR MEANS COMMITMENT

POINT-COUNTERPOINT

deploy software defined networks, but enterprises are a different animal.

"We're still very early in figuring out the use case definitions for software defined networks," said Michael Spanbauer, principal analyst with Current Analysis.

One of those use cases is certainly the enterprise data center, where virtualization has changed traffic patterns and added complexity. However, fabric solutions are aimed squarely at enterprise data centers with these problems, and they are shipping today.

"If I wanted to make a decision on an enterprise-class data center today at scale, I would have to go with one of the shipping fabric solutions," Spanbauer said. "But that doesn't mean you lock out OpenFlow as a solution down the path because all of [the fabric vendors] are active members of the [OpenFlow] working groups."

A FIRST STEP: INTEGRATING SDN WITH FABRIC WITH LEGACY NETWORKS

Many SDN proponents point to the ability of their technology to make multivendor networks programmable from a centralized control point using OpenFlow solutions that require OpenFlow-friendly switches and routers. However, companies like Nicira Networks, Big Switch Networks and VMware will use network tunneling protocols like VXLAN, STT and NVGRE to create a software overlay network that abstracts and virtualizes the physical network from the virtual servers at the access layer of the data center LAN. This capability could be helpful to enterprises that are incrementally adding modern data center fabrics to their infrastructure. SDN can abstract that pocket of modern fabric and the rest of the data center LAN.

"If you're using QFabric and other vendors in your data center, software defined networking could give you an overlay to that, which would allow you to manage your QFabric but then manage other infrastructure as well," said Juniper's Marcellin.

ARE DATA CENTER FABRICS AND SOFTWARE DEFINED NETWORKS COMPLEMENTARY?

Using SDN as an overlay to abstract management of multiple vendor footprints in your network is one thing, but SDNs and fabrics have much deeper complementary potential. Even in a data center where one vendor's network fabric serves the entire data center, SDN can deliver added value.

"The most intelligent and pragmatic solutions tend to look at the problem of [data center] network-

HOME

IDEA LAB

DATA CENTER NETWORK FABRICS VS. SDN

WITH NETWORK FABRICS, CHOOSING A VENDOR MEANS COMMITMENT

POINT-COUNTERPOINT

ing and break it into two pieces—the physical network piece and the dynamic network piece on top of that," said Internet Research Group's Christy.

The first piece—the physical network—is not trivial. "Software people don't understand how hard it is to build a physical network that is reliable and manageable. And one of the hardest problems is [how] to use all the connecting wires on demand for whatever needs to be done at that moment. That is a difficult problem, and it is the kind of thing that network vendors like Cisco, Juniper and Arista do very well," Christy said.

"Then there is the communication between the virtual machines, which increasingly you need to deal with in the realm of software. The problem you run into with traditional networking is when things become more dynamic and when more of the operation moves into the software at a higher level."

Big Switch Networks, an SDN and OpenFlow startup, has at least one joint customer with Juniper QFabric, according to Kyle Forster, co-founder and vice president of marketing at Big Switch. And he anticipates having many more, not just with QFabric but with data center fabrics from Cisco, Brocade and others.

"Software defined networking gives you the functionality, and a fabric gives you the bandwidth," Forster said. "If you want to place any virtual machine anywhere and you need to get subnets and VLANs fundamentally out of the way, you need a whole lot of management functionality. Most of that, and in my opinion, all of it is in the software defined networking camp."

A data center fabric can guarantee an enterprise equal bandwidth no matter where it chooses to place a virtual machine within a data center. As those VMs move around, the functionality of an SDN comes into play.

The joint QFabric customer, a cloud provider that caters to the healthcare industry, offers a multi-tenant environment for its customers. In that environment, it replicates a hospital's existing data center. Over time, the cloud provider dynamically evolves those environments, moving VMs and other resources around to reduce its own internal costs. However, as it alters its internal network to support the migration of those resources throughout its data center, the cloud provider wants that network to remain transparent to its customers, Forster said.

"That means that the bandwidth guarantees had to be there [from QFabric] as [the provider] moved physical servers or virtual machines around. And he really needed to be able to place any physical server or virtual machine anywhere in the

HOME

IDEA LAB

DATA CENTER NETWORK FABRICS VS. SDN

WITH NETWORK FABRICS, CHOOSING A VENDOR MEANS COMMITMENT

POINT-COUNTERPOINT

data center as he cost-reduced," Forster said. "He wanted to be able to do that without saying to the customer, 'Oh, hey, due to funky reasons that are more my problem than yours, I want to change your IP addresses.' That would be very distasteful to a customer relationship."

"This is an extreme case of [a customer] wanting to get the network complexity out of the way and the VLANs and subnets out of the way," Forster said. "We're giving him all the Layer 2/Layer 3 functionality, and QFabric is giving him all the bandwidth."

EVERY NETWORK FABRIC VENDOR IS BUILDING AN SDN STRATEGY

If you have any doubts on whether fabrics and SDNs are reconcilable, look at the fabric vendors. Cisco, Juniper, Brocade, Extreme Networks, Dell Force10 and Arista Networks are all articulating and evolving SDN visions. They are working to find points of intersections between the two categories of products.

"I can't speak for any vendor, but I suspect that every vendor is thinking about their migration strategy [to SDN and OpenFlow]," said Pitt of the Open Networking Foundation. "The general notion of fabrics is a really good idea. They are proprietary solutions that seek to solve some of the same problems that we are addressing with OpenFlow and software-defined networking in general. I think some of these original [fabrics] were by necessity proprietary, but I don't think there's much of a need to stay proprietary for much longer since we have a standardized approach that's now in the market." ■

If you have any doubts on whether fabrics and SDNs are reconcilable, look at how the fabric vendors are working to find points of intersections between the two product categories.

WITH NETWORK FABRICS, CHOOSING A VENDOR MEANS COMMITMENT

Can incompatible fabrics be good for data center networking?

BY MICHAEL MORISY

AS SPANNING TREE PROTOCOL (STP) shows its age in data center networking and Transparent Interconnection of Lots of Links (TRILL) and Shortest Path Bridging (SPB) crawl toward true standardization, network architects on the bleeding edge are left to choose between betting the server farm on a single, generally proprietary fabric vendor and waiting for the fabric wars to end.

With Extreme Networks, Brocade, and Cisco embracing TRILL, Avaya and Alcatel-Lucent supporting SPB and Juniper backing its own QFabric, it could be a long wait.

DESPITE UPDATES, STP SHOWS ITS AGE

First conceived of in 1985 by the then-giant Digital Equipment Corporation (DEC) and standardized in 1990 by IEEE, STP served as the default for routing meshed local area networks in a way that eliminated dangerous loops, corrected meandering paths and allowed redundancy.

But that flexibility came at a cost.

"Even with a set of enhancements in various forms over the course of many years, recovery in a spanning tree environment can take a long time," said Eric Hanselman, research director with the 451 Group. "It can take up to 30 seconds; the original spec took up to 90 seconds when there was a failure, and that's a long time for a network to be out."

STP also blocks redundant paths to prevent loops. This constrains bandwidth, which is increasingly in demand with the increase of east-

HOME

IDEA LAB

DATA CENTER NETWORK FABRICS VS. SDN

WITH NETWORK FABRICS, CHOOSING A VENDOR MEANS COMMITMENT

POINT-COUNTERPOINT

west traffic in highly virtualized data centers. This constrains bandwidth, which is increasingly in demand with the increase of east-west traffic in highly virtualized data centers. Driven by these shortcomings, the Internet Engineering Task Force (IETF) and the Institute of Electrical and Electronics Engineers (IEEE) went back to the drawing board and created two competing successor networking standards —TRILL and SPB—aimed at serving the needs of modern data centers.

While both standards have vocal industry and vendor support, the two approaches are generally incompatible; worse, many implementations of the same protocols are not even interoperable.

PRE-STANDARD IMPLEMENTATIONS OF STP ALTERNATIVES: VENDOR LOCK-IN?

"There's no such thing as TRILL interoperability," said Ethan Banks (CCIE #20655), a data center network engineer with an e-learning software company and a blogger at PacketAttack. "If you pick a fabric solution, you are really marrying yourself to that vendor."

That lack of interoperability—common with new standards—is not something the major vendors have even tried to hide.

Brocade's Jonathan Hudson recently wrote that his company and Cisco, two of the early TRILL leaders, "had no choice but to implement a 'pre-standard' TRILL variant." He stated that Cisco was TRILL compliant in the control plane but not the data plane, while Brocade was compliant in the data plane but not the control plane.

> "In the ideal world, fabric standards should have been done 10 years ago to give vendors time to build to the standards. But in the real world, data centers need to upgrade their infrastructure, and they need to control their costs."
>
> —DENISE SHIFFMAN
> *vice president, Juniper's Platform Systems Divisiony*

Both companies have said that as TRILL works its way toward final standardization, after years of committee wrangling, compatibility is on their roadmaps.

"All of our network switches which support FabricPath also work with TRILL," said Rajan Panchanathan, Cisco's director of project

HOME

IDEA LAB

DATA CENTER NETWORK FABRICS VS. SDN

WITH NETWORK FABRICS, CHOOSING A VENDOR MEANS COMMITMENT

POINT-COUNTERPOINT

management. As with many vendor-driven innovations of emerging standards, what started as proprietary Cisco extensions have often worked their way into the evolving standards, and the same would be true for many of the extensions Cisco has developed with its FabricPath implementation of TRILL, he said.

"Standards bodies take their time to get their heads around something and get consensus," he said. "Cisco is focused on solving our customers' problems. The standards bodies adopt most of or all of Cisco's innovations."

The compatibility picture is slightly better when it comes to SPB-based fabrics. For one, the technology was built to be backward compatible with modern data-center switches, easing the pain for those planning a staged upgrade.

But while SPB's backers have touted successful interoperability trials in a lab setting, the technology has failed to garner as much market adoption or interest as the decidedly incompatible TRILL.

"Shortest path bridging is a service provider play," said Banks. "There is some interoperability with the folks doing it, but it's not going to get a lot of traction in the enterprise space, I don't think."

Banks' Packet Pushers co-host, Greg Ferro, was a little more blunt in a blog post last year, decrying the SPB standard's capabilities and deriding its major backers, Avaya and Huawei, as bit players.

"I haven't had a single listener or person ever ask me about SPB," Ferro wrote. "While I have done some research, it's obviously not a topic of interest to anyone in my audience."

Juniper has opted out of both TRILL and SPB in favor of its proprietary QFabric product and is open about why it thinks those standards are overrated.

"There's an ideal world and there's a real world," said Denise Shiffman, vice president of Juniper's Platform Systems Division. "In the ideal world, fabric standards should have been done 10 years ago to give vendors time to build to the standards. But in the real world, data centers need to upgrade their infrastructure, and they need to control their costs."

That off-script innovation leaves network engineers like Banks in a bit of a lurch: Choosing a fabric vendor today means being comfortable with long-term commitment in a field where better offerings might spring up overnight—or solid-looking suitors might disappear.

"My advice is to know what you're getting into," Banks said. "Brocade has been rumored to be looking for a buyer for a long time. That sort of thing will factor into people's buying decisions."

HOME

IDEA LAB

DATA CENTER NETWORK FABRICS VS. SDN

WITH NETWORK FABRICS, CHOOSING A VENDOR MEANS COMMITMENT

POINT-COUNTERPOINT

FABRIC STANDARDS ARE SETTLING, BUT SHOULD YOU?

Despite all the turmoil, there are positive signs. TRILL and SPB are widely expected to be finalized as standards this year, and the two protocols' various backers have promised to work on compatibility, with plugfests already being planned (although not, notably, compatibility between the two protocols).

Nick Lippis, founder of Lippis Enterprises Inc. and organizer of one of those events, is bullish about the future.

"I think every enterprise is dealing with a huge challenge: data deluge," he said. "Everyone knows what's happening with mobile devices and what they've been driving in terms of data requirements and storage requirements."

That deluge is creating a sense of urgency around more efficient, redundant and available data centers. Independent confirmation of compatibility and performance will help move fabric adoption, Lippis said.

"The market should be heating up in the second part of this year and into 2013," he said.

For some companies with urgent data-center needs, 2013 might be too long to wait, Banks said, but that is likely to be the minority: bleeding-edge financial firms, service providers, the largest enterprises.

Banks looked at FabricPath as an option for his own company, but he found the cost too prohibitive at this point. "So we said we're not going to use that. We're going to use the virtual port solution, which isn't a fabric as such."

For now, passing on any fabric option, and keeping your eyes open towards the future, is generally an acceptable answer.

"The real question around fabrics is do they deliver the kind of benefits you need with the penalty of single-vendor support?" said Hanselman. "Most people using fabric today are using it on the scale where that single-vendor requirement isn't a penalty."

But even if a business can afford to wait and see what happens, it might be a good idea to start talking to vendors about what they can offer—while re-evaluating existing relationships.

"Everybody is doing a little of something different," said Banks. "You can just bring in the same sales rep you've had for five years, but it's a really smart time to bring in a vendor you've never talked to before and say, 'I'm due for a refresh. This is my problem. What do you have to solve this and why?'" ■

POINT-COUNTERPOINT: DO BOTH SDN AND DATA CENTER NETWORK FABRICS FAIL?

In this point-counterpoint feature, two network engineers-turned-bloggers—Ivan Pepelnjak of IOS Hints and Brad Casemore of Twilight in the Valley of Nerds—take opposing sides in the battle of network fabric vs. SDN. Read both, and see which side you agree with. BY RIVKA GEWIRTZ LITTLE

WHEN IT BECAME clear that virtualization and the cloud would strain the data center network—with new east-west traffic patterns, extreme application workloads and the need for flexibility and convergence—all of the major vendors began battling to prove they had the best solution for the problem. For each of them, that solution was a complex and costly data center fabric that promised flat, non-blocking, end-to-end transport between any node on the network. And then software-defined networking (SDN) came and shook things up.

Suddenly, proponents of SDN promised that the control plane of the network would be decoupled and centralized, making it possible to build a network of dumb devices that could be granularly controlled down to the individual traffic stream. This technology could be used to spin up virtual network instances on demand and treat compute, storage and networking merely as pools of flexible resources. With that kind of manageability and flexibility, who needed cumbersome traditional networking architectures? In fact, who needed data center fabrics?

In truth, neither data-center network fabrics nor SDN models are completely proven, so in which option (if any) should users invest? ■

HOME

IDEA LAB

DATA CENTER NETWORK FABRICS VS. SDN

WITH NETWORK FABRICS, CHOOSING A VENDOR MEANS COMMITMENT

POINT-COUNTERPOINT

POINT…

With SDN, no need for 'rococo' data center network fabrics

Once software-defined networking takes hold, there may be a need for Ethernet fabric, but it will be basic and exist only to serve the layers of SDN abstraction. **BY BRAD CASEMORE**

LET'S ASSUME, FOR sake of argument, that software-defined networking (SDN) fulfills its considerable promise and stakes claims not only to the data centers of large cloud-service providers, but also to those of large enterprises. Would there be any need for the Ethernet fabrics peddled by Cisco, Brocade, Juniper, and other established networking vendors?

There will still be a role for fabrics, but not for the relatively rococo fabrics the major vendors are positioning and selling today.

In certain key respects, SDN and network fabrics have much in common. Both emerged as potential network solutions to challenges posed by rampant data-center virtualization and cloud computing that brought new east-west traffic patterns and application workloads that strain the limits of longstanding architectures and technologies.

However, that's where the similarities end and the differences begin. The fabric offerings of leading networking vendors, typified by standard and non-standard approaches to Ethernet multipathing, represent a linear progression in networking's evolution. The networks are being flattened to varying degrees, and the shackles of Spanning Tree Protocol (STP) are being cast aside, but it's still very much business as usual for the vendors and their customers. The latter are continuing to see their vendors of choice present network infrastructure composed of distributed control planes, vertically integrated switches, and proprietary extensions to industry-standard protocols and technologies.

SDN takes a different tack. The purpose of SDN is to enable network virtualization and network programmability through the decoupling of the network control plane from the data-forwarding plane. In SDN, the network control and intelligence gets pushed up and out to server-based software (a new realm for vendors' proprietary value creation) and away from the underlying network hardware, which ideally will be fast, reliable, cheap, and relatively dumb.

The data centers that adopt and implement SDN will not need the proprietary fabric networks proffered by the major vendors, which for the most part are architecturally and philosophically at odds with the ONF's vision of virtualized infrastructure (compute, storage, networking) managed as integrated resource pools in thrall to application workloads.

But—and here's where things get interesting—they'll still need fabrics. Given the nature of SDN, these emerging fabrics are likely to be relatively simple. They'll be akin to physical networks that mirror the any-to-any connectivity of a chassis backplane. Eric Hanselman of the 451 Group says a network fabric is "a technology for extending chassis-like functionality across multiple physical systems." Likewise, Nicira CTO Martin Casado has written on his blog that "a fabric is a physical network which doesn't constrain workload placement....a physical network that operates much as a backplane does within a network chassis."

Both definitions seem apt for what SDN requires. In the long run, a base network fabric will emerge to serve the higher layers of SDN abstraction, but today's proprietary fabrics are unlikely to fit the bill.

Now, the question is, who will provide this fabric? With value and margins gravitating from the physical network to server-based controller and the applications that run on them, the major networking vendors will be disinclined to accept the role of subservient plumbers. They're used to richer margins and more control (pardon the pun) over their destinies. They will accept this new role grudgingly, if at all.

It's not likely that complex fabrics from proprietary vendors will win out.

Meanwhile, we'll see other vendors fill the void with relatively cheap hardware that meets the needs of the SDN community. Already we can see the ONF launching a concerted bid to persuade purveyors of merchant silicon (Broadcom, Marvell, Intel's Fulcrum) to deliver OpenFlow support in their chips. At the same time, ODMs have begun to deliver bare-bones switches to customers at cloud-service providers and other major data centers. They're following a trail blazed by server original design manufacturer's (ODMs) that similarly provided cheap and cheerful compute hardware for service-provider clientele looking to virtualized resource pools as a means of adapting and harnessing cloud computing.

It will take some time, but it's not likely that complex fabrics from proprietary vendors will win out. ■

COUNTERPOINT

Data Center Fabrics vs. SDN? Neither works

Without doubt, there is need for change in the data center, but network virtualization is a better answer than complex fabrics and SDN.

BY IVAN PEPELNJAK

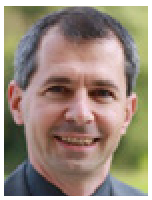

LIKE WITH EVERY new technology, proponents of OpenFlow and software-defined networking (SDN) keep promising to solve all the world's problems. There are very valid use cases for OpenFlow and SDN, but data center fabric, defined as the networking gear providing end-to-end transport within a data center environment, is not necessarily one of them.

It's important to first agree on the definition of SDN. While OpenFlow is well defined (although one has to wonder how open it is), SDN by itself is a meaningless term. After all, software has defined networking ever since IBM announced the 3705 Communications Controller in 1972. In fact, software has always controlled networks and large networks have always been at least partly monitored and configured by additional software. However, for the purpose of this discussion, we use the definition of SDN the Open Networking Foundation (ONF) proposed: an architecture where network control is decoupled from forwarding.

A number of similar architectures are already widely used in data center networks—and they're not that successful. Cisco's Virtual Switching System, Juniper's Virtual Chassis and QFabric Network Node, and HP's IRF all use central control-plane software that programs the forwarding tables in numerous switches. Yet all these architectures have a common trait—they don't scale. The most that networking vendors have managed to do thus far is to control up to 10 top-of-rack switches or eight core switches with a single control plane. NEC has managed to do slightly better with its ProgrammableFlow controllers. Also, if you take a closer look at Google's infamous G-scale network—the best example of production-grade OpenFlow we've seen so far—you'll see a similar picture. Google uses OpenFlow to control a small number of devices that are physically close to the OpenFlow controller. Managing a dynamic environment with very fast feedback loops and high rate of change from a central point simply doesn't work.

Don't get me wrong—I'm not saying today's data center networks are perfect. In fact, they are riddled with numerous scalability challenges, most of them stemming from the simplistic implementation of virtual

HOME

IDEA LAB

DATA CENTER NETWORK FABRICS VS. SDN

WITH NETWORK FABRICS, CHOOSING A VENDOR MEANS COMMITMENT

POINT-COUNTERPOINT

networking by most hypervisor vendors that use VLANs to create virtual networks. But instead of fixing the root cause of the problem—by moving the complexity to the hypervisor—the networking industry keeps proposing a series of kludges, with SDN/OpenFlow being just of them. Eventually they might manage to create flying pigs (and charge you extra for the jetpacks), but the hidden complexity of these solutions would be comparable to the complexity of the traditional voice networks. We only got to large-scale and low-cost voice when we stopped using voice switches and moved to peer-to-peer VoIP solutions like Skype.

SO WHAT MIGHT WORK?

Imagine a scenario (that large cloud providers like Amazon and Rackspace already use) where virtual networks get implemented with a MAC-over-IP (be it VXLAN, NVGRE or Nicira's STT), or an IP-over-IP solution, and storage engineers would use an IP-based solution (be it iSCSI, NFS or SMB 3.0). In this case, all the complexity is moved to the hypervisors, and it makes perfect sense to control that environment with OpenFlow because the controllers would be dealing with a very large number of independent uncoupled devices. This is what Nicira is doing with its Network Virtualization Platform. The data center network would have to provide just two services: end-to-end IP transport and lossless transport of specific traffic classes.

In a well-designed data center network that provides pure IP transport, you no longer need to change the switch configurations every time you add a new virtual network or a new server. The only time you have to change the network configuration is when you add new capacity, and even then the existing tools some data-center switch vendors offer allow you to automate the process.

To summarize: Once we get rid of the VLAN stupidities, and move the virtual networking problem to where it belongs (the hypervisors), we no longer need complex data center fabrics and complex tools to manage their complexity. Existing large-scale IP networks work just fine and won't benefit much from an SDN-like centralized control plane. On the other hand, having a decent provisioning tool for a large-scale IP network would be a huge benefit. We thus don't need SDN/OpenFlow in data center fabrics; what we need is a Puppet/Chef-like tool to build and deploy them efficiently. ■

David B. Jacobs of The Jacobs Group has more than 20 years of networking industry experience. He has managed leading-edge software development projects and consulted to Fortune 500 companies as well as software start-ups.

Lisa Phifer owns Core Competence Inc., a consulting firm specializing in business use of emerging Internet technologies. At Core Competence, she has advised large and small companies regarding security needs, product assessment and the use of emerging technologies and best practices.

David Davis is the author of the best-selling VMware vSphere video training library from Train Signal. He has written hundreds of virtualization articles on the Web, is a vExpert, VCP, VCAP-DCA and CCIE #9369, with more than 18 years of enterprise IT experience.

Brad Casemore provides consulting and advisory services to technology companies. He helps identify and pursue market opportunities in fast-growing and competitive areas. He also helps companies identify and build strategic business partnerships that support their objectives.

Ivan Pepljnak, CCIE No. 1354, is a 25-year veteran of the networking industry. He has more than 10 years of experience in designing, installing, troubleshooting and operating large service provider and enterprise WAN and LAN networks and is chief technology advisor at NIL Data Communications, focusing on advanced IP-based networks and Web technologies.

TechTarget

Networking Evolution Ezine is a SearchNetworking.com e-publication.

Rivka Gewirtz Little
Executive Editor

Shamus McGillicuddy
Director of News and Features

Michael Morisy
Feature Writer

Kara Gattine
Senior Managing Editor

Linda Koury
Director of Online Design

Kate Gerwig
Editorial Director

Peter Larkin
Associate Publisher
plarkin@techtarget.com

TechTarget
275 Grove Street, Newton, MA 02466
www.techtarget.com