

**April 1, 2014**  
**Jeremy Li**

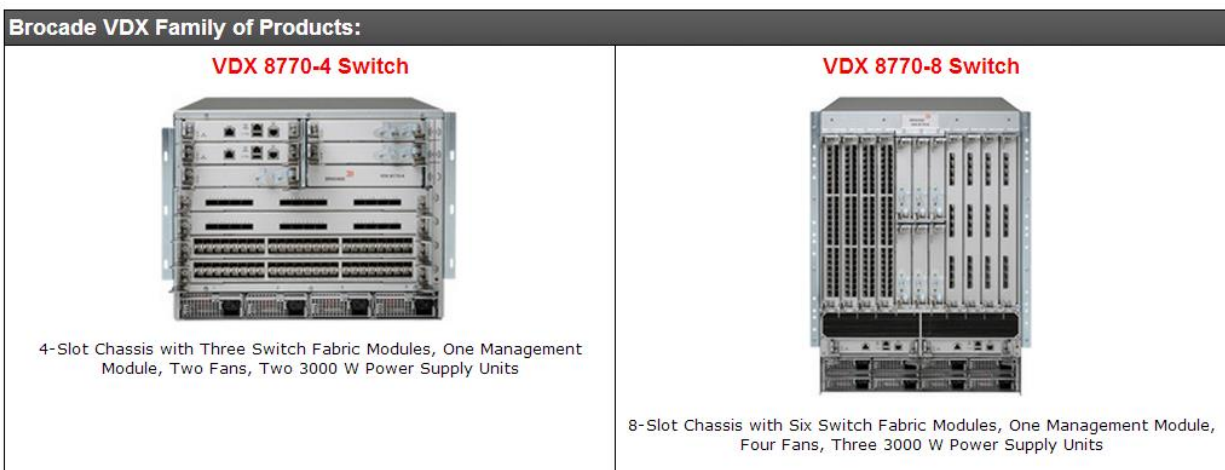
A Brocade Virtual Compute Block (VCB) is reference architecture for integrated infrastructure, designed to cost-effectively and quickly scale virtualization and cloud computing.

The trend is to switch from Ethernet box to an Ethernet fabric that is an Ethernet construct, a repository, where multiple physical switches are interconnect with each other in addition to having software running on hardware to form elastic, automated and simpler switch network. Most simpler switches may not have an embedded & complicated software so that they will concentrate on accepting and passing packets quickly and let Controllers manage switches. This idea led to create a Software Defined Network (SDN).

Brocade has Virtual Cluster Switching (VCS), which was developed to provide a lossless, low latency, deterministic multi-path Ethernet network.

Brocade claims that it offers linear scalability from hundreds to thousands of virtual machines (VMs) with dynamic allocation of compute and storage resources under its new Data Center (DC) Virtualization architecture.

If the scalability described by Brocade referenced above is true, Brocade's VCS Ethernet Fabric that is comprised of VXD 8770 switches, as illustrated in the screenshot below (see [Brocade's VDX datasheet](#) for details) will lower the TCO dramatically since most enterprises will not rely on third party tools such as [VMTurbo](#) and etc.



In summary, Ethernet fabric is the reasons for “agility, simplicity, scalability, efficiency as well as automation”.

Now, we have two major schools that promote Ethernet fabric: Cisco's and the rest of remaining switching companies, although many vendors have their own proprietary Ethernet fabric varieties.

VPLS agreed with Jeremy how to move legacy three-tier architecture to Ethernet Fabric architecture by creating a new data center without disrupting its current legacy data center (DC) in the future because Jeremy mentioned the fact "since we are still relying on Cisco Spanning Tree Protocol (SPT) that is instability, inefficiency and inability to scale due to its single-rooted tree and single-path nature, meaning "do not support virtualized traffic", a new DC will address this problem, instead of upgrading our current DC to avoid any disruption."

Q1: Can you tell me more about the major difference between Cisco's approach and Brocade's approach?

A1: Brocade offers 32 VDXs switches in single fabric with 8,000 ports capability to address the East-to-West traffic in modern data center.

Its VDX switches are based on a field-programmable gate array (FPGA), an ASIC chipset that is an integrated circuit designed to be configured by a customer or a designer to enhance its capability when a newer technology is available. As a result, Brocade can offer its product longevity, meaning avoid a fork-lift upgrade, while Cisco customers who purchased Nexus 7000 core switches in 2013 must either replace their newly acquired Nexus 7000 products with Cisco's flagship product, often refer to as Nexus 9000 in order to support Cisco's new generation and prevailing Software Defined Network (SDN), Cisco now refers it as Application Centric Infrastructure (ACL).

- Brocade offers a 30% lower cost in its initial acquisition in comparison with Cisco's offering, meaning a 30% lower cost in CAPEX
- Brocade usually reduces the OPEX cost by 50% in comparison with Cisco's product because Brocade VDX products rely on automation, while Cisco Nexus 7000 or 9000 product requires a manual configuration, often, with CLI.
- Both Brocade and Cisco support OpenFlow version 1.3 in its selected model, while HP offers OpenFlow support across all its products.

**Note:** Brocade also offers a "pay-as-you-go" model in its VDX family to allow customers to pay for only the required capacity.

- VDX 8770 supports: 1/10/40GbE
- VDX 6740 supports: 1/10/40GbE
- VDX 6710 supports: 1/10GbE

VDX 8770 is in the same level with Cisco's Nexus 7000, but is designed to support Software-Defined Networking (SDN) implementations across data, control, and management planes, while Nexus 7000 does not support Cisco newer technology such as ACL.

Cisco also moves towards **leaf- spine** architecture, meaning two-tier architecture in its new data center (DC) solution in addition to supporting L2 ECMP pre-standard TRILL fabrics and L3 ECMP fabrics in its Ethernet Fabric architecture.

Q2: Are Cisco Nexus 9000 and ACL proprietary? Can you provide a comparison sheet regarding VDX 8770 vs. Nexus 9000?

A2: VPLS will not be able provide the answer in the follow-up email.

Note: See [Cisco reveals OpenFlow SDN killer](#) for details

Briefly discussed the following topics:

- LAG, Link Aggregation Group - similar to Ether Channel to aggregate more ports into a single virtual port
- MC-LAG, Multi-Chassis Link Aggregation Group - aggregate more physical switches into a single virtual switch
- TRILL, Transparent Interconnection of Lots of Links – Pre-standard stage at this time

**Note:** Juniper keeps away from TRILL for the purpose in favor of its own proprietary technology, which is based on its Virtual Chassis scheme that combines physical switches into a single logical device in order to utilize its [QFabric](#) architecture for seeking flatten (often one tier network) and simplify data center networks

- SPB, Shortest Path Bridging – mostly used by Alcatel and Lucent Telco
- Other vendor's proprietary protocol (e.g., Juniper's proprietary L2 ECMP)

Q3: Can you kindly tell me the SDN protocols listed below:

- VXLAN – Virtual Extensible LAN
- STT – Stateless Transport Tunneling
- NVGRE - Network Virtualization Generic Routing Encapsulation
- VMware 5.5 introduced a new feature called Virtual Distributed Switch (vDS) that might be incompatible with Cisco's approach.
- VMware NSX gateway (SDN) vs. Cisco's approach?
- OpenFlow - is a [communications protocol](#) that gives access to the [forwarding plane](#) of a [network switch](#) or [router](#) over the network (Source: Wiki)

A3: The SDN is virtually replacing vDS because Brocade's VDX physical switches can be configured as a single logic switch.

Q4: As a VM-aware network, Ethernet fabric should have [federation](#) and integration with VM management software such as vCenter or System Center to support workload mobility within or across data centers with full automation for network policy and resources.

Does Brocade VDX have a capability referenced above?

A4: Yes

Q5: How many SDN architectures do we have besides VMware SDN and Cisco's ACL?

A5: See [Cisco reveals OpenFlow SDN killer](#) for details.

Q6: What is SDN device-based or overlay-based or hybrid-based SDN?

A6: VMware SDN is based on software, while others third party vendors use hardware based SDN controller.

Recommendation: First to use a freebie from VMware SDN. If it will not meet the needs in an enterprise, then, look for a hardware-based SDN.

**Note:** VMware might offer licensing free if an enterprise is in the VMware Enterprise Plus licensing agreement.

Typically, a \$12,000 VMware licensing fee is for one dual-socket ESXi Host.

Q7: Tell me more about cloud orchestration platform or software (e.g., OpenStack)

A7: Audio Conferencing time was expired.

Q8: Some vendors provide northbound API programmable capabilities for their Ethernet fabrics as well. Tell me more about it?

A8: That ACI policy model submission will also include a northbound Group Policy API from Cisco's ACI work, the company said.

Cisco and industry partners are all contributing to this release, and OpenDaylight will ensure an "open" approach and transparency to Cisco's policy-based technology, Cisco says. Still, observers note a not-so-subtle Cisco bend in the direction OpenDaylight is taking.

Q9: Since southbound proprietary SDN protocol is used to control fabric devices, is Northbound API or protocol used for network management?

A9: [Cisco reveals OpenFlow SDN killer](#) – now refer to as OpFlex that will become Cisco's favored southbound protocol. Below is a quote from the above link:

"In essence, Cisco is re-inventing the OpenFlow wheel, proposing a new protocol where one already exists, though its objective is different. OpFlex will be to Cisco's Application Centric Infrastructure (ACI) [programmable networking architecture](#) what OpenFlow is to SDNs."

Q10: Should enterprises rely on “OpenFlow” as a standard to avoid a proprietary protocol OpFlex?

A10: That’s a tough question. Read a quote from the link below for details:  
[Cisco reveals OpenFlow SDN killer](#)

“Cisco is lining up ACI partners in physical and virtual switching, and Layer 4-7 network services, to support OpFlex and write to its APIs. Cisco is proposing it as a standard in the IETF and, with partners IBM, Plexxi and Midokura, submitting OpFlex to the OpenDaylight open source SDN project – which it [co-founded with IBM](#) -- for an ACI-compatible policy model OpenDaylight plans to offer in its upcoming “Helium” release.”

One day, the newly released Cisco proprietary protocol OpFlex might become a standard in the future because an IT heavyweight IBM is behind the current proprietary protocol OpFlex.

Q11: Does Brocade VCS powered by VDX switches separate *Control Plane* and *Data plane*?

A11: Read a quote below from a good article titled “[Cisco reveals OpenFlow SDN killer](#)”

“OpenFlow could not have served this role. It's meant to serve a model where the control and data planes are decoupled, and in which the switch's intelligence is limited to packet forwarding.”

**Note:** According to Dave, if SDN is enabled, both *Control Plane* and *Data plane* will be decoupled.

Q12: What kind of Data Center Architecture in Microsoft and Azure (Public Cloud)?

A12: Stay tuned...

In summary, Kevin Walsh from Brocade states:

Here are our keys:

1. Our Fabric are self-healing, self-load balancing, VM interactive and aware of all Vmotions etc. They are dynamically adjustable (expand or contract) with very limited interaction and are unmatched in east-west traffic performance. Speed, performance, configuration, we crush it.
2. Our Fabrics are incredibly easy to manage. Our fabric can be managed as 1 device. Simple. To add a Brocade switch to our Fabric, 2 commands. Our competitor, literally over 100. Brocade’s interaction with VMWare means any movement of VMs requires no interaction of admins, we auto sense, auto correct

and auto populate all movements in the fabric. Automated, simple, robust and fully meshed, as a fabric is.

3. We traditionally are anywhere for 20-50% less expensive on the front end and over the life of the equipment can be even greater savings.

We are 100% confident in our ability to deliver.

I don't believe Nexus 9000 does FabricPath at all. If it does, that is new and I would be really careful because frequently with Cisco, even within models, certain features don't play with certain line cards and just trying to keep up with it is insanity. (e.g. you can't mix M1 and F1 and F2 cards in the same N7K chassis and have FabricPath work...just to name 1 out of dozens of very quirky series and chassis configurations issues with cards and features sets)

### **Recommended Reading:**

1. [Fabric wars: Cisco vs. Brocade vs. Juniper](#)
2. **DATA CENTER SHIFT:** [HP accuses Cisco of diverting data center standard](#)

### **Acknowledgement**

Thank John Minnix, his management team and Kevin Walsh for their time to present Brocade Ethernet Fabric Overview via audio conferencing.