The Economic Advantages of Open Networking and Customer Choice

Choice of industry-leading network operating systems on Dell Open Networking switches

DELL EMC

Table of Contents

1. The enterprise IT challenge

Enterprise organizations are dealing with highly dynamic computing, storage and networking needs at a greater capacity than ever before. The primary factors driving the industry are mobility, cloud computing, big data, and the Internet of Things (IOT). However, the static nature of the traditional network is not designed to address the contemporary data center network and computing needs of evolving businesses. These new dynamics have brought about a shift, introducing virtualized network traffic patterns, resource pooling, cloud, and the need for server-like automation. Additionally, this shift in dynamics has led to modern data center network environments which are designed to simplify and automate physical switch deployments.

The current networking industry looks much like the proprietary computing market of 40 years ago, where a small number of industry players produced vertically-integrated solutions with software and hardware that were single-sourced. This practice led to a proprietary industry culture where innovations were stifled and operations were fragile, complicated and expensive.

> "[The] new era of networking is defined by disaggregation, software & virtualization, and open source."
> — Guru Parulkar, Executive Director, Open Networking Research Center

While IT organizations have attempted to meet these new dynamics using the legacy network infrastructures mentioned above, this approach has resulted in increasingly complex deployments that make network management and maintenance even more burdensome and costly.

According to a recent report from Delphi Technology, many organizations are moving away from legacy platforms to more modern vendor supplied solutions. This trend is based on increasing support and maintenance problems with the current legacy systems. In some

[1] http://www.riskandinsurance.com/wp-content/uploads/2015/09/Delphi_Opportunity-Cost-of-Retaining-Legacy-Systems.pdf

cases the legacy system's technology is a few generations behind resulting in inefficient processes and management. Another example includes technical staff moving on leaving an organization struggling to maintain a system that many IT staff are unfamiliar with. A final example is the astronomical costs of some support agreements when a vendor is charging more for the effort required to upkeep aging technology.

A key contributor to IT's inability to move quickly is the condition of the existing network infrastructure. While other systems within the data center are accelerating and adopting new technologies, the status you of networking has remained stagnant. Many of today's proprietary networking vendor solutions maintain a vertical "black box" approach where hardware and software are closely joined aimed at locking in customers.

"The only way to stay ahead of this growth curve is software."
— John Donovan, Chief Strategy Officer and Group President,
AT&T Technology and Operations

Fortunately, innovations in networking are constantly being introduced and, as a result, solutions to these challenges are available. The next section will take a closer look at some data center networking advances that will enable enterprise organizations to improve their network infrastructures while also reducing costs.

2. The emergence of spine-leaf fabrics in the data center

In traditional data center networking, engineers faced a fundamental challenge: in order to increase scale and redundancy, they were required to add multiple tiers, devices and physical links. Because of these extra links, Spanning Tree Protocols (STP) had to be configured in order to ensure a loop-free topology between all of the data center switches. These protocols were required for stability, but also caused more complicated and expensive architectures with several downsides: unused bandwidth, slow convergence, and increased management complexity.

Figure 1 – Legacy spanning tree examples

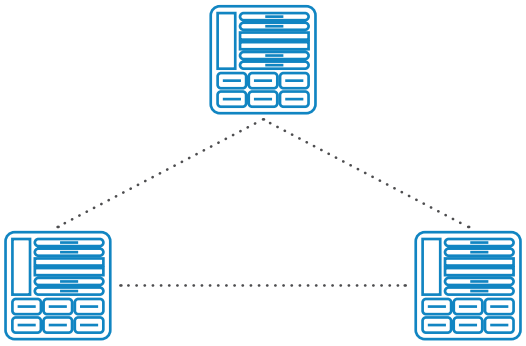

Simple 3-Switch STP Example

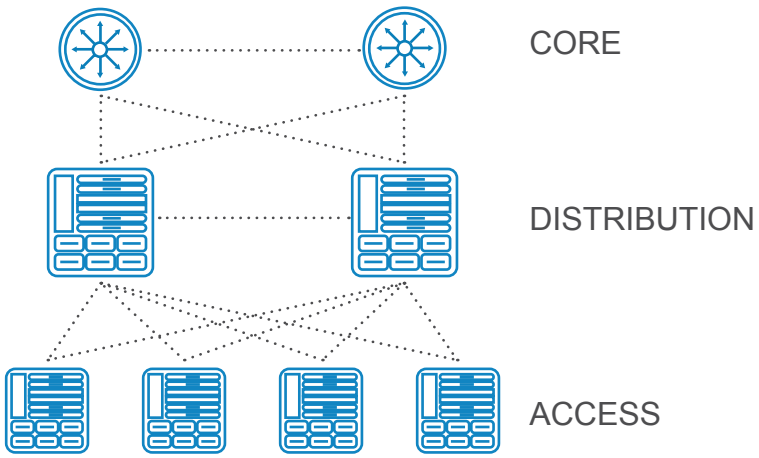

Industry vendors began developing protocols and solutions to address these downsides. The first attempt at solving some of the STP issues leveraged new Layer-2 multi-path protocols such as Transparent Interconnection of Lots of Links (TRILL) [RFC6325]. TRILL resolved several of the issues with STP, but suffered from limited industry adoption. The main concerns with adopting TRILL included vendor proprietary extensions, complex configuration and management, and sub-optimal traffic handling, along with the high cost of procuring hardware capable of supporting TRILL.

Due to these issues with TRILL, industry quickly shifted away from the “tiered” topologies to Clos-based fabric topologies, or spine-leaf fabrics. These data center fabrics are optimized for today’s modern virtualized data centers, which are experiencing increasingly heavy east-west traffic patterns based on the new dynamics discussed earlier. This fabric design flattens the physical topology, provides a predictable switch-to-switch latency, and largely removes the risk of network loops.

Figure 2 – Spine-leaf topology

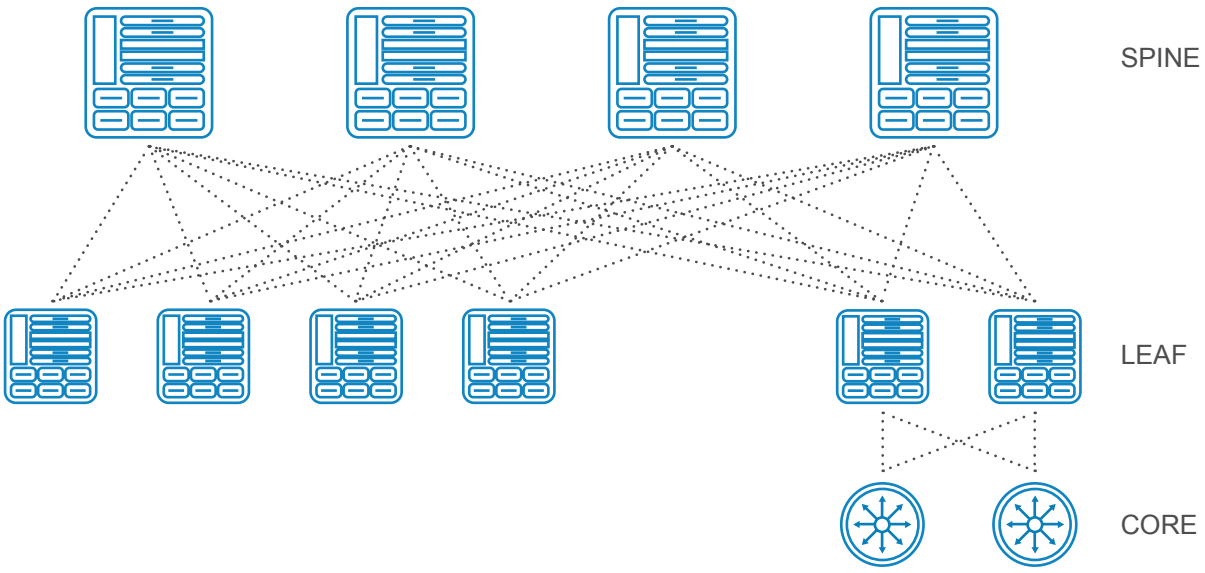

These spine-leaf fabrics, originally popularized by hyperscale data-center operators such as Google, Facebook, and Amazon, are an increasingly popular design for enterprise organizations due to their elegant scale-out properties, resiliency and attractive cost dynamics:

- **Scale-out:** In a spine-leaf fabric, every leaf switch connects to every spine switch, with every link being active all the time.
- **Resiliency:** As there are a large number of parallel active links, several links can fail before connectivity across the system as a whole is materially impacted.
- **Cost:** Spine-leaf fabrics are typically constructed of smaller (1-2U) industry-standard fixed-form switches. Compared to proprietary and expensive chassis-based designs, the smaller, fixed-form factor switches are less expensive, consume less rack space per port and require less cooling. They provide a low latency non-blocking solution, allow for oversubscription and can be swapped out easily.

Dell was an early pioneer in data center network fabrics. Our Active Fabric solutions employ Clos spine-leaf architectures with our Open Networking fixed form-factor switches. We offer a family of switches which maximizes customer choice, flexibility and innovation at any scale. The next section of this paper will discuss how Dell is shaking up the market with our industry-leading open networking solutions.

3. Dell's market changing Open Networking initiative

As proprietary mainframes and mini-computers from companies like IBM and Digital were replaced with commodity-based x86-based servers, the computer industry experienced tremendous innovation, technology velocity and cost reduction. The decoupling, or disaggregation, of the operating system (OS) from the computing hardware allowed consumers the freedom to select the best and most efficient hardware and software components at each layer, while also customizing solutions to best meet their strategic business objectives and cost requirements.

Unfortunately, despite the adoption of industry-standard data center networking protocols such as Ethernet, the networking industry remains stuck in the 1980's proprietary era, where switches from several major networking vendors are comprised of locked-in hardware/ software stacks that are not interchangeable with other vendor offerings.

Dell's vision follows our earlier success as a pioneer in disaggregating computing hardware and software, offering an open data center network ecosystem in which organizations can pick and choose from a selection of innovative, industry-standard network applications, network operating systems, and network hardware. This open approach enables businesses to simplify their data center networking through complete disaggregation, including disaggregating the operating system from the hardware, disaggregating the virtual network from the physical network, and disaggregating the control plane from the data plane.

Figure 3 – Dell Open Networking

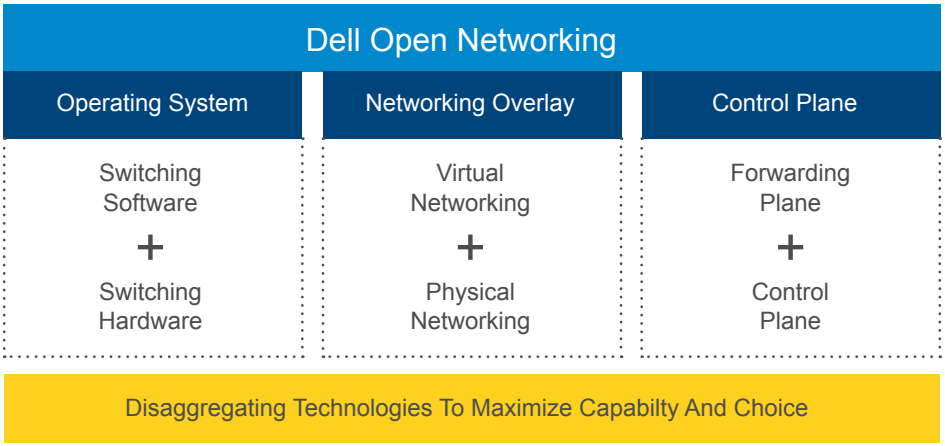

As an early global end-to-end technology company supporting a choice of third-party network operating systems, Dell enables organizations to tailor their data center networks to specific applications within their organization. Our disaggregated-networking model has disrupted the traditional proprietary networking paradigm: we give organizations an open alternative for their unique needs with a growing open networking ecosystem of solutions and partners.

Dell truly believes that our open network solutions stimulate rapid innovation by helping our customers achieve unprecedented levels of flexibility and efficiency. These solutions also help minimize the time and effort required to design, provision and manage networks; enable IT managers to leverage open-source tools; and provide expertise to help reduce costly engineering overhead.

In the next section, we'll take a closer look at some management challenges when deploying fabrics in the data center and how Dell offers both our own and 3rd party partner solutions to alleviate these challenges.

4. Simplified management of data center fabrics

Since their popularization five years ago, one of the key challenges with data center spine-leaf fabrics has been their complex management. When using legacy network management technologies, even a simple spine-leaf fabric involves complex control protocols. As a result, outside of the hyperscale data centers, this exciting fabric network topology has seen slow adoption.

Dell's Active Fabric Manager (AFM) is a direct answer to the rapidly changing dynamics that challenge today's enterprise. AFM offers compelling benefits to enable enterprise IT to deliver powerful new capabilities that drive the business bottom line. When running Dell's Operating System 9 software, Dell's AFM empowers enterprises to immediately and easily deliver a highly-automated SDN-enabled Ethernet fabric. By leveraging the latest innovations in SDN and network programmability, AFM can automate the design, deployment and day-to-day operations of data center fabrics. This automation reduces fabric deployment time by up to 86% compared to manual configurations while eliminating costly configuration errors that account for the majority of outages today.

Figure 4 – Dell Active Fabric Manager

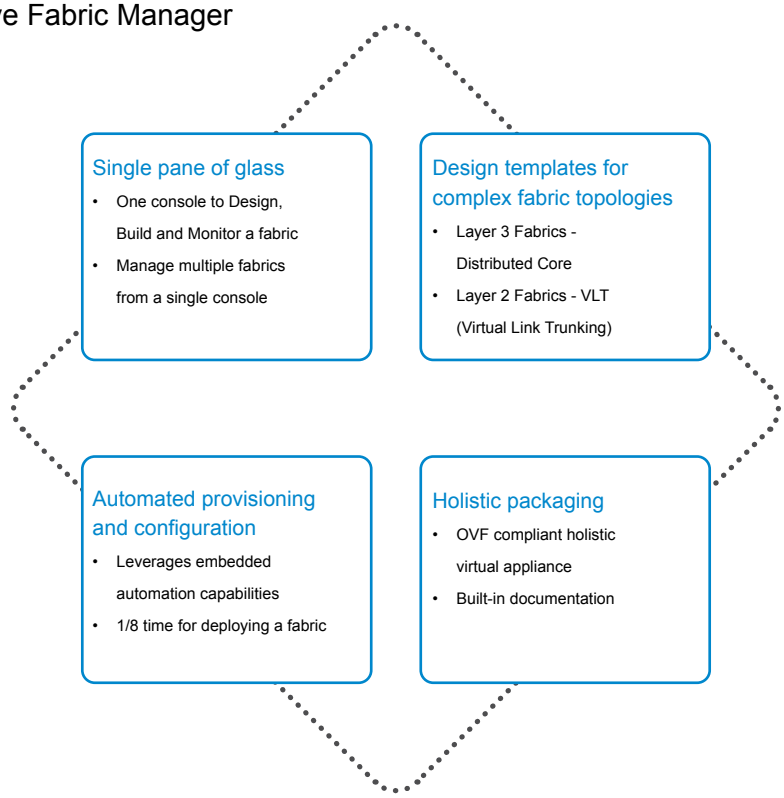

While enterprise IT needs to adopt new technologies quickly, complexity and significant training requirements make adopting and mastering the latest technologies a daunting challenge. AFM's powerful software tools enable IT staff to masterfully design and deploy next generation Layer 2 and Layer 3 fabrics within a matter of minutes. Unlike traditional network management tools that focus primarily on performance monitoring of various network elements, AFM has been purpose-built to leverage SDN technology and deliver a new operational model where automation is the baseline.

For customers choosing to take advantage of Dell Open Networking solutions running alternative OS's, there are two new innovative partner approaches for managing spine-leaf Clos fabrics with different software solutions that use exactly the same hardware.

- **Big Cloud Fabric from Big Switch Networks** allows data-center operators to manage spine-leaf fabrics as if the fabric is one big switch, using Software Defined Networking (SDN) controllers that act like supervisors, spine switches that appear to the user like a chassis backplane, and leaf switches that appear to the end user like chassis line cards. Complete with productized integrations for vCenter, OpenStack, and other cloud orchestration systems, this approach to spine-leaf fabric management is highly automated yet intuitive for traditional networking professionals.

- **Cumulus Linux from Cumulus Networks** provides data center operators the entirety of the native Linux experience and vibrant ecosystem of applications. Leveraging the existing open source and commercial Linux applications and toolsets enables operators to manage the network in the same seamless workflow as their Linux server. With examples on how to incorporate popular Linux management and DevOps tools like Puppet, Chef, Ansible, CFEngine, or SaltStack, this approach to spine-leaf fabric management is intended for IT shops looking to fold networking into their server management tool chains.

The following sections list a few examples where management innovation, made possible by Dell Open Networking, provides the advantages of spine-leaf fabric design deployed by leading technology companies and service providers to enterprise organizations on a much broader scale.

Example 1: Big Cloud Fabric versus Cisco FabricPath: simplified management

Figure 5 – Comparing Dell/Big Cloud Fabric to Cisco FabricPath

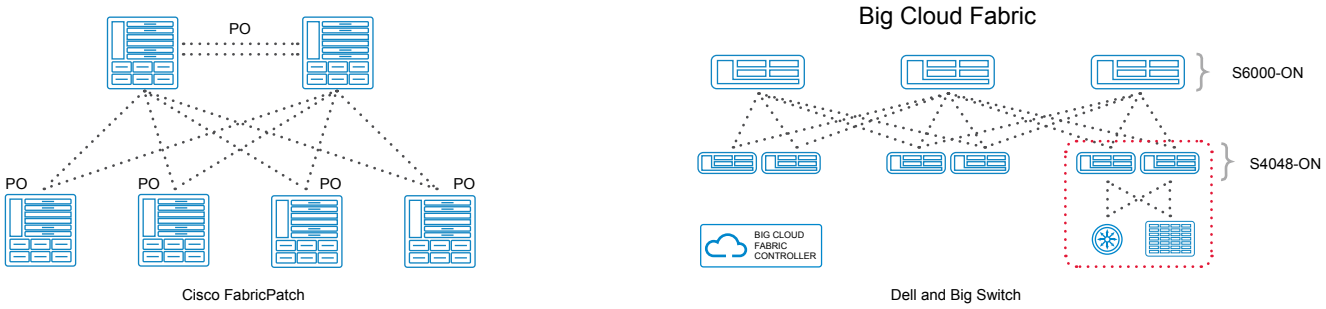

Using a typical customer-use case of a 16-Top-of-Rack deployment, we compare Cisco's Nexus series switches running FabricPath and Cisco's version of the IETF TRILL standard to Dell Open Networking partner Big Switch Networks Big Cloud Fabric solution, a modern SDN fabric running on Dell Open Networking switches.

Table 1 – Big Cloud Fabric comparison chart

	Big Cloud Fabric	FabricPath
Configuration management points	1	44
Ubiquitous L2	Yes	Yes
L3 boundary	Everywhere	Only at top layer
Device image upgrade	Single hitless	Per device
Multi-tenancy	Yes	No
RESTful API	Yes	No
End-to-end analytics	Yes	No
VMware integration	Yes	No
OpenStack integration	Yes	No
Rack units	20	44
Power (Watts)	5233	19020
BTU/HR	17858	57040

The chart above quickly highlights some of the advantages (both from a CAPEX and OPEX perspective) a newer SDN fabric has over an older TRILL-based deployment. These advantages help organizations to move forward and adopt a modern technology deployment rather than continuing to support legacy infrastructure.

Savings

- 55% on space
- 72% on power
- 69% on cooling
- 46% on networking cost
- Single point of management (44 to 1)

Example 2: Big Cloud Fabric versus Cisco: better automation

Beyond the obvious advantages of having a single point of management with SDN controllers on a spine-leaf fabric, Big Cloud Fabric (BCF) software with Dell Open Networking hardware also ships with fully productized software integrations for both OpenStack and VMware environments.

Many of the frequent workflows running across data center fabrics such as application deployment, or error-prone workflows such as Link Aggregation Group (LAG) formation, are fully automated in this model. Basic network troubleshooting can be accomplished via vCenter consoles or OpenStack Horizon dashboards before going to more advanced networking-specific tools using the BCF Graphical User Interface (GUI), RESTful Application Program Interface (API), or industry standard Command Line Interface (CLI). This rich integration provides concrete and tangible results that improve the lives of data center admins and provide immediate value.

Table 2 – Big Cloud Fabric workflow advantages

Workflow	8 Rack OpenStack pod	16 Rack VMware pod
	Cisco N7k/3K vs Dell and Big Switch	Cisco N7k/3K vs Dell and Big Switch
5 Year CapEx (HW/SW)	Target 50% less	Target 50% less
Application deployment	16x faster	16x faster
Initial set-Up	8x faster	8x faster
LAG/fabric formation	12x faster	18x faster
Software upgrade	20x faster	30x faster
Connectivity troubleshooting time	12x faster	12x faster
Pod expansion	12x faster	N/A

Along with various enhancements to support migration to BCF in brownfield VMware environments, a new BCF plug-in for vCenter 6.0 demonstrates a much tighter integration for the network and virtualization admins.

Table 3 – BCF integration with VMware

BCF automation	vSphere environment	vSphere + NSX environment
Host detection	Automatic	Automatic
Host LAG formation	Automatic	Automatic
L2 network creation	Automatic (for port groups)	Automatic (VTEP, vMotion, and storage port groups)
VM learning	Automatic	Automatic
Network policy migration during vMotion	Automatic	Automatic
BCF visibility and troubleshooting	**vSphere environment**	**vSphere + NSX environment**
ESXi host inventory	Yes	Yes
VM-to-host association	Yes	Yes
Fabric analytics for virtualization	VM-level visibility (VMname, vMotions, Hostname…)	VM & NSX level Visibility (VMname, Logical Switch, VXLAN/VNI…)
Troubleshooting physical underlay	Yes (VM-to-VM)	Yes (VTEP-t0-VTEP)

Example 3: Cumulus Linux and automation harmony

Dell Open Networking partner Cumulus Networks takes a different approach to solving the deployment and management challenges of the spine-leaf fabric. By leveraging off-the-shelf Linux-based configuration management tools, monitoring software, and DevOps frameworks that data center System Administrators already use, helps reduce the months and weeks of deploying network configurations to seconds. Additional advantages include:

- Ease of deployment
- Enforce Network Configuration Policy
- Common set of tools across the data center
- Ideal for OpenStack, Hadoop as well as traditional network designs

- Interoperable with existing network infrastructure
- No controller needed, uses traditional routing protocols network engineers understand
- Massive scale, battle tested network designs and award-winning support

Figure 6 – Cumulus Networks spine-leaf design

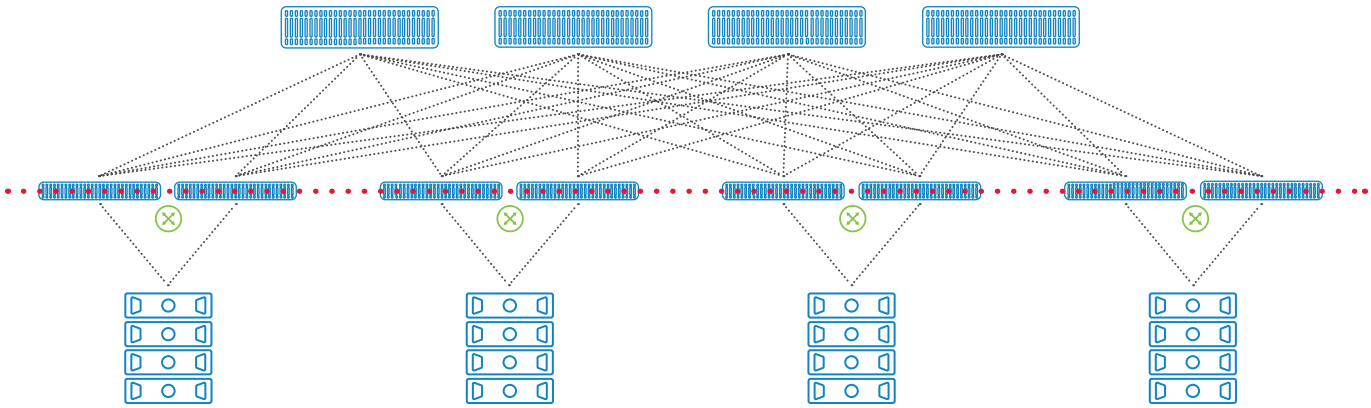

Using the same requirements from above, the following chart highlights the advantages a Cumulus Networks and Dell Open Networking solution would deliver.

Table 4 – Cumulus Linux comparison chart

	Cumulus Linux	FabricPath
Configuration management points	1	44
Language extensibility (Bash, Perl, Python, Ruby)	Yes	No
Unnumbered interfaces (BGP and OSPF)	Yes	No
Native DevOps tools support (Ansible, Chef, Puppet, Salt)	Yes	No
Native Linux troubleshooting (scamper, mtr, iperf, ping, traceroute)	Yes	No

Native Linux tools for monitoring (Nagios, Ganglia, influxdb, collectd)	Yes	No
Full interoperability with other vendors	Yes	No
Variety of methods to configure Host to ToR(VLAN Trunks, MLAG VXLAN, RoH, Anycast)	Yes	No
Multi-tenancy	Yes	No
Rack units	20	44
Power (Watts)	5233	19020
BTU/HR	17858	57040

Savings

- 55% on space
- 72% on cower
- 69% on cooling
- 57% on networking cost

Example 4: Data Center Shared Vision

Dell, VMware, and Cumulus Networks share the same vision that data centers can be more agile and flexible when defined in software, and not hardware-based. These three companies are working together to deliver solutions to help enterprises and service providers rapidly provision physical and virtual networks, then deploy new applications within minutes. These solutions significantly simplify IT operations and improve IT response time.

Figure 7 – Dell VMware and Cumulus solution

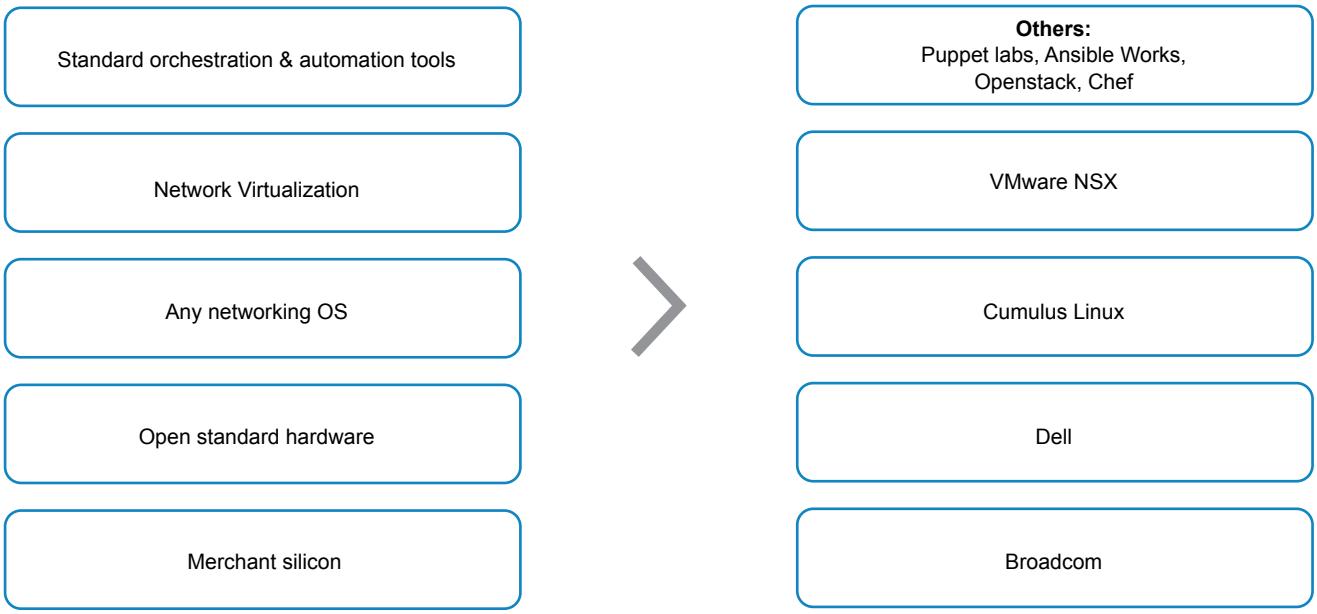

Bringing together VMware software with Cumulus Linux on Dell networking switches enables:

- Unified visibility and control of end-to-end for virtual and physical infrastructure
- Integrated solution stack with the intelligence distributed in software without vendor lock-in
- Simple consumption model for customers across solutions
- Feature velocity with software innovation due to the untethering of hardware and software
- Lower Total Cost of Ownership (TCO) with the increase in operational efficiencies and CapEx gains

5. Conclusion: the benefits of an Open Networking ecosystem

Traditional IT deployments leveraging proprietary networking vendor locked-in solutions, obsolete software and hardware, and antiquated certification processes are expensive, inefficient, and inflexible. Continuation of these processes will result in slow and expensive solutions that impede an agency's ability to improve services and quickly respond to ever-changing needs and regulations.

The earlier use cases provide just a few examples of what Dell Open Networking and disaggregation can enable. Further opening of the data center network for additional software will allow for many additional possibilities such as:

- Tap-less network analytics
- Network function service chaining
- Controller-less IP fabrics
- Container optimized networking
- Hybrid/private cloud orchestration
- Automatic threat and Decision Disk Operating System (DDOS) response

Similar to how virtualization changed the computer industry, Dell's disaggregation, SDN, and open networking architectures have the potential to change the networking landscape forever. Dell, along with its ecosystem of partners, is assisting enterprise organizations to seamlessly migrate to a far simpler, open and innovative model of business already embraced in the private sector. Doing so will enable enterprise IT departments the ability to select state-of-the-art solutions at every layer, resulting in:

- Lower operational and capital expenditures
- Improved network performance
- Improved service delivery times
- The ability to remain flexible to adopt new innovations

To learn more about our open networking solutions, details of the comparisons data or methodology used, please contact your local Dell representative or visit Dell.com/networking

January 2017