Extreme®
Connect Beyond the Network

White Paper

An Introduction to VCS Fabric Technology

Future-proof your network with VCS Fabric technology, a cost-effective solution purpose-built for cloud environments.

VCS Fabric technology, which provides advanced Ethernet fabric capabilities, enables you to transition gracefully to elastic, highly automated, cloud-optimized networks.

Industry discussions of server virtualization, SDN, and cloud computing abound, for one simple reason: agility. Every year, CIOs are asked to support more—often much more, in terms of data and traffic—with less, and more quickly. In addition, a number of regulations have arisen in the last decade, forcing a stronger focus on data security and control. These business realities, along with improvements in WAN/LAN performance, have contributed to a retrenchment of IT operations and investments back into the data center.

VCS® Fabric technology is designed to meet these challenges by enabling next-generation virtual data center and private cloud computing initiatives. VCS fabrics, delivered on VDX® switches, were designed specifically to meet the needs of cloud environments and scale-out data center architectures being driven by an explosion of information, the heightened importance of data, and new application architectures. VCS fabrics provide the highest levels of network automation, efficiency, resilience, and investment protection. This paper reviews key data center network challenges and how VCS Fabric technology addresses them.

Ethernet Fabrics

Compared to classic hierarchical Ethernet architectures, Ethernet fabrics provide higher levels of performance, utilization, availability, and simplicity. They have the following characteristics at a minimum:

- **Flatter**: Ethernet fabrics eliminate the need for STP, while still being completely interoperable with existing Ethernet networks.
- **Flexibility**: Can be architected in any topology to best meet the needs of any variety of workloads.
- **Resiliency**: Multiple "least cost" paths are used for high performance and high reliability.
- **Elasticity**: Easily scales up and down at need.

More advanced Ethernet fabrics have innovated more aggressively in the control and management spheres.

- They are self-forming and function as a single logical entity, in which all switches automatically know about each other and all connected physical and logical devices.
- Management can then be domain-based rather than device-based and defined by policy rather than repetitive procedures.
- These features, along with virtualization-specific enhancements, make it easier to explicitly address the challenges of VM automation within the network, thereby facilitating better IT automation.

Introduction

In survey after industry survey, CIOs consistently say they are driven to simultaneously improve business processes while reducing IT costs. These business concerns require data center operators to deploy new applications quickly and efficiently, provide fast and reliable "around-the-clock" access to information, and meet or exceed stringent service levels with no downtime. Addressing these business needs is a set of technology enablers, including high-density, multi-core servers, server and storage virtualization, and a move towards service orchestration and cloud computing.

However, there are important differences in how the "new" data center must be architected, since server-to-server transactions now exceed data center–client exchanges by a significant margin, according to Gartner. IT domains that had developed fairly independently in the client-server era have become much more interdependent, and the ability to coordinate resource allocation synchronously across both physical and virtual servers, storage, and networks is increasingly necessary to guarantee application availability.

Modern distributed mission-critical applications must be resilient, scalable, and highly efficient, placing profound new pressures on the networks that serve them. With today's application architectures, performance itself is now highly dependent on the efficiency and utilization of the network. At the same time, organizations need the ability to scale out their networks on demand—without adding complexity—to respond to immediate business challenges.

With regard to the data center network, Gartner has estimated that:

- "Data center network architects need to plan for at least a doubling in bandwidth from physical servers during the next two to four years" due to increasing VM density.[1]
- Driven by more live VM migrations, as well as growth in composite applications, "new traffic patterns will appear arbitrary and even chaotic, with fluctuations that can be 90 times higher than the traffic peaks experienced by most data centers today."[2]

As a result, Gartner predicted in 2011 that more than 80 percent of traffic in the data center network would be between servers by 2014."[3]

In addition, the rise of cloud computing has created a newly pressing need for more infrastructure agility—meaning elasticity, architectural flexibility, manageability, and simplified coordination across IT domains—as a building block for service automation.

Recognizing these trends early on, Brocade brought its 15-year heritage in fabrics to bear on building new networks for Ethernet in the virtualized data center. Brocade released the first Ethernet fabric technology, VCS Fabric technology, in December 2010. Through the Brocade acquisition, Extreme now has the largest Ethernet fabric install base in the industry.

[1] "Is Your Data Center Edge Network Ready for Virtualization?" (Munch, 4/21/11, ID G00210549).
[2] "Your Data Center Network Is Heading for Traffic Chaos" (Munch, 4/27/11, ID G00210674).
[3] Ibid.

VCS Fabric Technology

When Brocade set out to build an Ethernet fabric technology, it did so with several customer challenges in mind:

- Network barriers to improving application performance
- Guaranteed application availability and performance in the face of network disruptions or bottlenecks
- The need to elastically scale networks to support variable growth of traffic and other scale-out architectures, like scale-out NAS, with existing personnel
- The desire to quickly orchestrate customized network resources as part of a cloud-based service

Accordingly, VCS Fabric technology is built upon three core design principles:

1. It should "just work," with minimal human intervention.
2. It must efficiently support non-stop networking in a modern, highly virtualized environment.
3. It is built for the needs of cloud environments, with support for a variety of approaches to multi-tenancy, policy customization, and cloud orchestration.

According to IDC:

"Over the next decade, the number of servers (virtual and physical) worldwide will grow by a factor of 10, the amount of information managed by enterprise datacenters will grow by a factor of 50…Meanwhile, the number of IT professionals in the world will grow by less than a factor of 1.5."

"Extracting Value from Chaos" June 2011

Simple and Automated—"It Just Works"

Traditionally, network management has been labor-intensive, accounting for a very high percentage of overall data center network TCO. Administrators must possess extensive knowledge of a wide range of commands and protocols. Each device must be configured individually, and in many networks, switches in different tiers run different operating systems. In a classic three-tier architecture, increasing network scale has a multiplicative effect on the number of switches that must be purchased and deployed, as well as on the level of effort required to do so.

This mode of operation was manageable while traffic growth and virtualization adoption both remained at low to moderate levels. However, all industry data points to a rapid explosion in the data that is generated, exchanged, and stored, while most observers believe staff growth will remain relatively flat.

The need to automate IT operations in general, and networking in particular, is abundantly clear. This does not mean simply relying on broad IT orchestration and management tools, or even SDN, to help manage existing devices at a higher level than before. If these approaches are used instead of rather than in addition to fundamental network reform, they will simply mask or relocate complexity rather than actually reducing it, which may even complicate troubleshooting. Instead, the complexity of every element of data center infrastructure must be reduced and ease of alignment between IT domains improved, so that the higher-level IT management stack can provide better visibility and control of data center operations as a whole.

TRILL

TRILL is a new standard for delivering Link Layer (Layer 2) multipathing and multi-hop routing. Unlike STP, with TRILL the shortest paths through the network are active, and traffic is automatically distributed across the equal-cost paths.

Network Automation

A VCS fabric is designed to be managed as a single "logical chassis," so that each new switch inherits the configuration of the fabric, allowing the network to scale out with ease. This eliminates the need to manually configure and manage each switch, simplifying management and reducing operational costs. VCS Logical Chassis simplifies ongoing operations, with single-command firmware updates, centralized monitoring, and troubleshooting. For more information about VCS Logical Chassis, read An Overview of VCS Logical Chassis White Paper.

VCS fabrics are self-forming and self-aggregating, making them very elastic and enabling real-time scaling. The fabric is automatically aware of all devices (servers, switches, storage devices, and appliances) within its domain. The VCS fabric control plane improves upon the TRILL control plane, in that it provides automatic assignment of RBridge IDs, automatic resolution of duplicate IDs, and automatic ISL formation and topology discovery. The VCS Logical Chassis further simplifies deployment by automatically sharing global configuration information to all incoming switches. Physical and virtual servers can be located upon connection without the fabric requiring manual reconfiguration. In addition, VCS fabrics provide fabric-level APIs and extensions to OpenStack Neutron ML2 to efficiently orchestrate both physical and logical networking resources as part of Virtual Machine deployment to support multitiered application topologies.

Zero-Touch Virtualization Support

Today, most IT organizations say that approximately 50 percent of their workloads are virtualized. However, one of the major barriers to greater virtualization adoption to date has been the difficulty of consistently aligning network connectivity and services with VMs, another factor in the relatively slow adoption of VM mobility.

Clearly, the network has an important role to play here. Seventy-one percent of IT organizations are running virtualization technologies from at least two different vendors, according to ESG. VCS Fabric technology is hypervisor-agnostic, providing basic VM alignment capabilities that match the reality of most data centers. AMPP and VM-aware network automation features enable customers to fully align virtual server and network infrastructure resources and realize the full benefits of server virtualization.

VM-aware network automation provides secure connectivity and full visibility to virtualized resources with dynamic learning and activation of port profiles. In VMware environments, the VCS fabric communicates directly with VMware vCenter to eliminate manual configuration of port profiles. The VCS fabric also supports VM mobility across VCS fabrics within a data center, while providing protection against VM MAC spoofing. Additional VMware vCenter integration with Brocade Network Advisor provides another layer of intelligence to network administrators. (See Ease of Use of VCS Fabric Technology Augmented by Brocade Network Advisor At-A-Glance for more details.)

True "plug-and-play" fabric capabilities that automatically align physical and virtual resources are critical for cloud deployments, where services that rely on a virtualized infrastructure may be turned on or off in real time. Only VCS Fabric technology delivers this capability today.

Network Efficiency

Resilience is a requirement in modern data centers with clustered applications and demanding compute SLAs. In developing VCS Fabric technology, Brocade naturally implemented this core characteristic in its Ethernet fabric design in an innovative fashion.

Resilience in the Ethernet world has traditionally meant massive over-deployment of resources. This is because in traditional Ethernet networks running STP, only 50 percent of the links are active; the remaining links (shown as dotted lines in Figure 1) act as backups in case the primary connection fails.

The traditional three-tier Ethernet network shown in Figure 1 was designed for "north-south" traffic between physical servers running monolithic applications and clients. However, such a design is also inflexible, as well as difficult and expensive to scale. In environments with heavy "east-west" traffic within server clusters, this traditional design also imposes significant latency burdens.

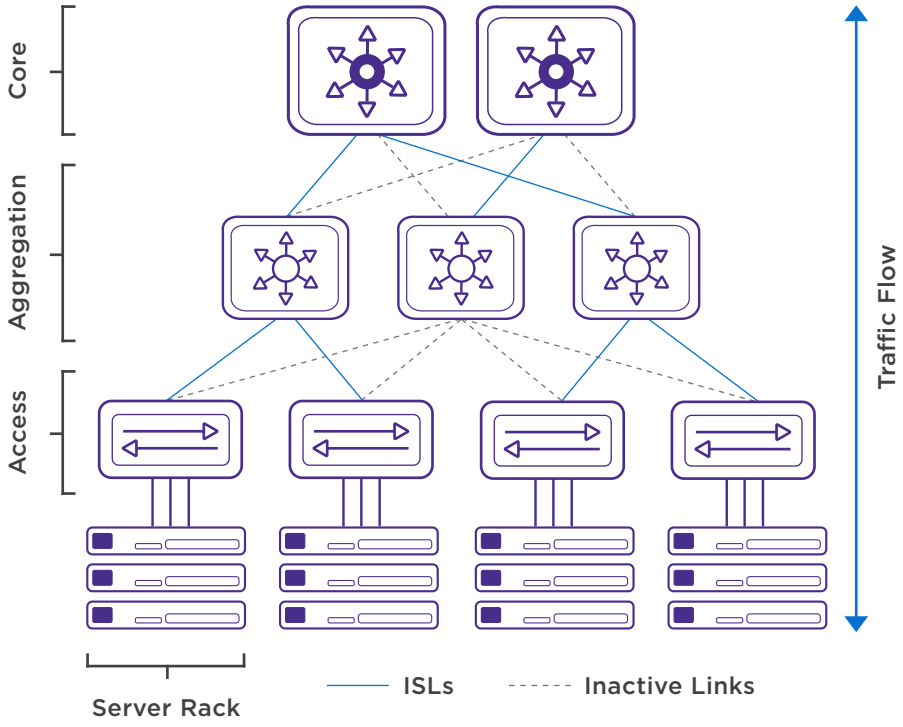

Figure 1 Classic Ethernet data center network.

In addition, it is considered a best practice to constrain the scalability of a traditional Ethernet network, because the larger the network, the longer the network takes to reconverge when a failure or other network change occurs. However, constraining the size of the network affects not just scale but flexibility, particularly in virtualized environments. VM mobility must occur within the bounds of a Layer 2 network; if the network remains small, the sphere of VM mobility also remains small.

VCS Fabric technology is architected for efficiency at every level. At the highest level, it is based on a logically flat, access/aggregation-centric design rather than a three-tier, core-based architecture, for several reasons:

- Increasing VM density means increasing I/O at the access point. By eliminating STP within the access layer, VCS Fabric technology immediately doubles the available links and bandwidth to the servers, without multiplying the capital and facilities costs of twice as many switches.
- VCS fabrics can also be implemented in the aggregation tier, either above existing traditional access switches or as part of a multi-fabric design. This approach can provide significant scalability for the attached hierarchical network while sheltering core switches from considerable traffic.
- Core switches are the most expensive part of a data center network. By implementing Ethernet fabrics at the access and/or aggregation tiers rather than in the core, VCS Fabric technology offers a low cost of adoption while preserving existing core investments.

VCS Fabric technology does not use STP; rather, it is TRILL-based. Therefore, VCS fabrics allow for active-active connections originating from the server, enabling much more rapid failover and convergence in the fabric. VCS fabrics are also self-monitoring; the Fabric Watch feature tracks health at the switch component level, while the Logical Chassis construct enables fabric-wide configuration and monitoring through a single point of management. In the event of an outage, links can be added or modified quickly and non-disruptively. This self-healing fabric approach doubles the utilization of the entire network, while improving resilience. It also allows IT architects to confidently increase the size of their Ethernet networks, which helps make VM mobility much more feasible.

VCS Fabric technology enables very elastic Layer 2 and Layer 3 domains with more effective and responsive link utilization and improved resilience at all layers of the network. VCS fabrics provide four or more active Layer 3 gateways and highly efficient load balancing at Layer 3, on top of Layer 2 ECMP and automatic ISL Trunking, a patented technology. These capabilities provide efficient inter-VLAN traffic optimization, since all VDX switches can handle Layer 3 traffic. Layer 3 traffic is kept as low as possible in the network topology, which minimizes hops, accelerates server-to-server traffic, and increases application performance.

In addition, each switch enabled with VCS Fabric technology switch has a full, fabric-wide view of Layer 2 address tables, as well as QoS, security, and VLAN-related information. Since the entire fabric shares the same virtual IP and MAC addresses, default gateways do not need to change when a VM moves, irrespective of the physical location of the VM within the VCS fabric. (See the Inter-Switch Link [ISL] Trunking At-a-Glance and Multilayer Multipathing in VCS Fabrics White Paper for more details.)

VM-aware network automation also provides secure connectivity and full visibility to virtualized resources with dynamic learning and activation of port profiles. AMPP, implemented in a hypervisor-agnostic manner, enables seamless VM migration, since the VCS fabric is aware of port profiles and automatically tracks them as they move. This helps ensure consistency across all fabric elements and minimizes downtime due to human error.

Finally, unlike other approaches to Ethernet fabric, which may require a central or "master" device, VCS fabrics are masterless and can be designed in full-mesh, partial mesh, leaf-spine, and various other topologies. With VCS Fabric technology, different end-to-end subscription ratios can be created or fine-tuned as application demands change over time, and the network can quickly and easily be reconfigured with minimal disruption to ongoing operations.

Built for Cloud

Public and private cloud providers need to deploy and support componentized, virtualized workloads quickly, securely, and scalably, on a per-tenant basis. Traditional VLANs can be used for this purpose up to a point, but limitations on VLAN ID scale and the rapid growth in numbers of tenants restrict their usefulness in larger data centers. The VCS Virtual Fabric feature of VCS Fabric technology is designed to address the scalability restrictions of traditional VLANs used for multi-tenant segmentation. The VCS Virtual Fabric capability provides scalable native multi-tenant support for both physical and virtual application deployments. Managed centrally through VCS Logical Chassis, VCS Virtual Fabrics simplify and accelerate application deployment and ensure policy consistency for each tenant, regardless of how application components are distributed across the data center. VXLAN and VRF Lite are other options for network segmentation. To learn more, read the Multi-Tenancy Options in VCS Fabrics White Paper.

Increasingly, SDN is seen as means to achieve control and customization of specific flows within multi-tenant data centers, and ultimately to improve service quality and agility. VCS fabrics themselves provide a first level of abstraction with their distributed control plane and the management abstraction paradigm of Logical Chassis. VCS fabrics also readily support overlay networks deployed in a variety of configurations. VDX Switches provide hardware support for overlay protocols of all types, ensuring synchronized visibility of physical and logical flows between overlay endpoints. The VDX 6740 and VDX 6940 Switches and VCS Fabric technology enable a VXLAN gateway across any point of the fabric.

Programmatic control of specific flows is enabled through industry initiatives, including Open Networking Foundation (ONF) and Open Daylight Project, as well as integration with VMware NSX and support for third-party SDN controllers. The VCS Logical Chassis helps simplify and scale SDN deployments by presenting a single interface to the SDN controller and offloading policy execution within the fabric. With support for OpenFlow 1.3, an industry-standard SDN communications protocol, it allows operators to address complex network behavior, optimize performance, and leverage a richer set of capabilities.

Similarly, VCS Fabric technology provides a single logical view to a cluster of physical switches to higher-level orchestration frameworks such as OpenStack—by leveraging the service-level abstractions defined in Logical Chassis mode via the fabric-level REST API. This offloads the orchestration of individual network elements from the orchestration platform, allowing it to achieve larger-scale, yet less complex deployments. Read the VCS Fabrics: The Foundation for Software-Defined Networks White Paper for more details. With support for Puppet and Python scripting, offers choice and more effective configuration management.

In addition to being optimized for virtualization, VCS Fabric technology also supports a range of storage access options (for example, iSCSI, NAS, and FCoE, as well as native Fibre Channel via several of the VDX Switches), as well as innovative features such as auto-QoS for NAS, so your network can evolve naturally with your storage strategy.

In short, VCS Fabric technology offers a cost-effective approach for simplifying and automating your network operations and provides a strong foundation for cloud adoption.

[4] Brocade VCS Fabric technology has no design limit on gateways. Please refer to the Brocade Network OS Release Notes for the current number of gateways supported.

Summary

VCS Fabric technology allows IT organizations to create efficient data center networks that "just work." Ethernet fabric architectures built on VCS Fabric technology share information across nodes and network layers, greatly simplifying network architecture and management, removing topological constraints, and reducing operational overhead. VCS Fabric technology offers unmatched VM awareness and automation versus traditional architectures and competitive fabric solutions, as well as features for supporting future "software-defined" networks across data, control, and management planes.

Only VCS Fabric technology, backed by a heritage of proven fabric innovations, delivers IT agility and assures reliability, with a cost-effective point of entry to allow you to transition gracefully to elastic, highly automated, mission-critical networks in cloud data centers.

VCS Fabric technology is embedded in the VDX switch portfolio. VDX switches are available today to enable IT organizations to build Ethernet fabrics to support cloud-optimized networking and greater enterprise agility.

VDX Switches

As IT organizations look for better ways to build clouds and virtualized data centers, they are turning to high-performance networking solutions that increase flexibility through leading-edge technologies. VDX switches are specifically designed to improve network utilization, maximize application availability, increase scalability, and dramatically simplify network architecture in virtualized data centers. VDX switches with VCS Fabric technology enable organizations to revolutionize the design of data center networks and provide an intelligent foundation for cloud computing.

VDX switches provide a flexible choice for building an Ethernet fabric, using two switches initially and scaling to add additional switches as demand increases. IT can mix and match 1 GbE and 10 GbE fixed VDX switches in an access-layer fabric using the VDX 6740 Switch for cost-effective 1 GbE and 10 GbE connectivity. The VDX 6740 also offers 40 GbE uplinks. The VDX 6940 offers the markets highest 10 GbE and 40 GbE ports density in a fixed form factor, making it an ideal choice for spine. The modular VDX 8770 Switch can be used for 1 GbE or 10 GbE connectivity in port-dense middle/end-of-row access fabrics and in aggregation-tier fabrics with 40 GbE and 100 GbE uplinks.

The portfolio of VDX switches provides Ethernet storage connectivity for FCoE, iSCSI, and NAS storage solutions within a single product family. IT organizations can protect their Fibre Channel investment by connecting Fibre Channel SANs to Ethernet fabrics with the VDX 6740.

VDX 6740 VDX 6940 VDX 8770-4 VDX 8770-8

Figure 2 VDX Family of Switches

Glossary

The following is a list of acronyms used in this document and their definitions.

Term	Definition
AMPP	Automatic Migration of Port Profiles
API	Application Programming Interface
ECMP	Equal-Cost Multipathing
FCoE	Fibre Channel over Ethernet
iSCSI	Internet Small Computer Systems Interface
250	NMS-250
ISL	Inter-Switch Link
NAS	Network-Attached Storage
QoS	Quality of Service
REST	Representational State Transfer
SAN	Storage Area Network
SDN	Software-Defined Network
500	NMS-500
SLA	Service Level Agreement
STP	Spanning Tree Protocol
TCO	Total Cost of Ownership
TRILL	Transparent Interconnection of Lots of Links
VLAN	Virtual LAN
VM	Virtual Machine