

Brocade One Data Center Cloud-Optimized Networks

Brocade® One™ represents a smooth transition to a world where information and applications reside anywhere in the cloud. To help customers make that transition, the Brocade One™ Data Center strategy simplifies network architectures and maximizes solution flexibility and choice. Unlike competitive approaches, Brocade One unlocks the full potential of the network without dictating a specific design or components.

Cloud Computing

Today's business priorities require data centers that deploy new applications quickly and efficiently, provide fast and reliable "around-the-clock" data access, meet or exceed stringent service levels with zero downtime, and do all of this while maximizing investments by reducing costs. In short, IT must move at the speed of business to capitalize on new opportunities and respond to increasing global competition. For this reason, many customers are creating private clouds in their data centers while leveraging public cloud services where appropriate.

Cloud computing represents a new business model in which IT requires service-based computing where shared resources are connected to applications, and cost is based on resource consumption, not asset purchase. Resource sharing in the cloud relies on server virtualization. This has triggered a dramatic shift in computing technology, so that applications are no longer "owned" by the server; instead, resources are assigned to virtual machines, which become the basic unit of computing. Virtual machines can move when workload change demands it and physical servers require upgrades and maintenance. And, they can move between private and public cloud data centers.

Cloud computing is a new economic model for data centers. It relies on virtualization of physical resources (compute, storage, and network) to reduce capital and operating costs, while reducing application deployment time from months to minutes. Clouds run in the data center either privately owned (aka private clouds) or owned by a third-party provider (aka public clouds). In either case, the goals are the same: reduce capital costs, reduce operating cost by efficiently sharing resources to optimize utilization, and reduce application provisioning time to minutes.

But, as data centers move to cloud-based architectures to increase enterprise flexibility and agility, the adoption of server virtualization is adding complexity to the network. Data centers are going to evolve so that they can deliver any application, anywhere, at any time. They are going to need a cloud-optimized network that seamlessly connects both the physical and virtual components, resulting in a flatter, simpler, intelligent, scalable, and efficient network.

Impact of Cloud Computing on the Data Center Network

Cloud computing is a new model for the data center. How applications consume resources and how applications connect to clients is virtual rather than physical. The use of the word “cloud” comes from IP networks, in which the physical network and data path are hidden from applications. Applications needed only an IP address to make a connection with other applications. An IP address virtualizes the network, hiding all the physical details from applications.

The idea of hiding physical infrastructure from an application is at the core of cloud computing. All physical components are hidden from applications. Before cloud computing, a physical server was used by a single application, which wasted capital resources, real estate, power and cooling, and human resources. In the cloud, physical servers are pooled into a server cluster and virtual machines are exposed to applications. The virtual machine hosts an application, or in the case of Web service applications, an application component. Resources consumed by applications (compute, storage, network bandwidth, and so on) are still provided by the physical infrastructure, but virtual machines acts as a “service layer” between the physical resources and the application workload that consumes them.

Application resource allocation and Service-Level Agreement (SLA) management is now tied to the virtual machine, not the physical server. The virtual machine becomes the basic unit of work in the data center, since physical resources are allocated, managed, and tracked by virtual machines, and the application inside consumes the container’s resources. Virtual machines must logically connect to several types of physical networks, for example, server clusters, data center management, client access, and storage networks.

When an application requires more resources than are allocated to its virtual machine, the virtual machine can be non-disruptively moved to another physical server in the cluster with more available resources, maintaining the SLA while improving resource utilization. When servers require maintenance or upgrades, all virtual machines using that server are migrated to other servers in the cluster so applications are “always on” even during planned maintenance windows. However, virtual machine mobility requires that all networks used by the virtual machine be aware of the migration to keep network policies, port configurations, Virtual LAN (VLAN) membership, and so on consistent. Network policies must be portable so they can follow virtual machines. In addition, Quality of Service (QoS) and traffic management services must be available end-to-end to maintain the application SLA; and internal network connections, or inter-switch links, must be able to handle abrupt changes in traffic patterns to avoid congestion *without manual reconfiguration*.

Client access connections to an application also have to follow the virtual machine non-disruptively, while continuing to be load balanced efficiently and secured across the server farms hosting virtual machines. Clearly, cloud computing places many new demands on traditional data center networks.

Networking Challenges

In the cloud, network resources are consumed by virtual machines, not physical servers. Network management has to extend beyond the physical to the virtual layer. Monitoring has to be more intelligent so it can correlate virtual machine traffic with physical network components and resources. QoS, bandwidth allocation and traffic shaping are important in network configuration so that application SLA requirements are met. Network configuration becomes more complex, requiring more administrative oversight.

Traffic patterns are much different for server clusters than for client/server connections. In server clusters, traffic between racks (“east-west”) can be much greater. Classic Ethernet networks support only “north-south” traffic flow, which adds latency to server cluster traffic moving between racks since that traffic has to move up and down through several network tiers. Server clusters benefit from a flatter Layer 2 network that scales across multiple racks so that east-west traffic flows use the shortest path to reduce latency and multiple paths for high utilization.

Since virtual machines can move, the network needs to know how to map them (and their hosted application workload) to its network policies. Virtual machines are assigned a persistent MAC address, which is associated with the appropriate network policies, as if it were a physical server. But, the policy has to follow the MAC address as it moves from one physical network port to another.

Workloads such as storage traffic require lossless, deterministic networks. Fibre Channel over Ethernet (FCoE) relies on Ethernet—which means that for lossless packet forwarding, Ethernet switches must support Data Center Bridging (DCB) extensions. Originally iSCSI relied on TCP for lossless packet forwarding, but it can also benefit from lossless Ethernet to avoid congestion more efficiently at Layer 2. Top of rack switches must also provide 10 Gigabit Ethernet (GbE) connections to the server using mezzanine cards with 10 GbE Converged Network Adapters (CNAs) and the growing use of 10 GbE LAN on Motherboard (LOM) connections. With CNAs, client IP traffic and storage traffic use the same physical link. But, bandwidth and QoS have to be carefully managed so that critical storage traffic does not experience congestion. For many customers, integration of FCoE traffic with existing Fibre Channel SANs is essential for cost-effective storage management, disaster recovery, and asset protection.

The manual configuration of multiple switches and ports with common parameters, as required in today’s network, does not scale, is not agile, and costs too much. With the addition of virtual machines and their logical connection to the physical network, converged IP and storage traffic on the same wire, and virtual machine mobility, network management must become agile and much easier to scale.

What Is a Cloud-Optimized Network?

In the data center, the cloud-optimized network connects virtual servers and their application workloads to other virtual servers and to virtual LUNs in storage resource pools, while maintaining secure, load-balanced client access connections to the Web, applications, and databases deployed on virtual machines. A high-level model is shown in Figure 1.

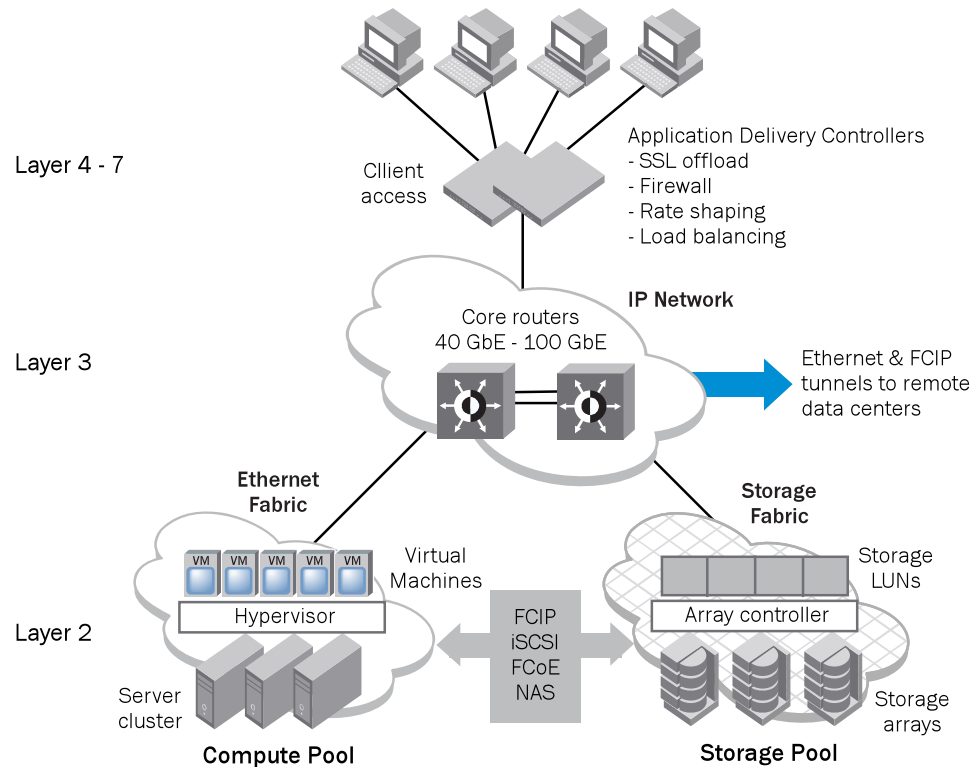


Figure 1. Model for a cloud-optimized network

The cloud-optimized network connects to virtualized resource pools and physical clients, servers, and storage using a single logical network. *Ultimately, a cloud-optimized network can consolidate multiple physical networks onto a common physical infrastructure.*

In Figure 1, an Ethernet fabric and storage fabric are implemented at Layer 2 to flatten the network, reducing capital and operating costs. Like storage fabrics before them, Ethernet fabrics are self-aggregating, scale efficiently, and are lossless and deterministic. In Ethernet fabrics, all switches are aware of all devices so virtual machine mobility does not require manual reconfiguration of the network. Finally, the fabric is extensible between data centers via core routers and Ethernet tunnels in the IP network.

Virtual machines with their applications will now move across a server cluster “stretched” between the private and public cloud data centers. Cluster traffic runs through the Ethernet tunnel. Storage traffic can also be tunneled over IP using the industry-standard Fibre Channel over IP (FCIP) protocol so that application data can be quickly replicated between public and private cloud data centers.

In cloud computing, the compute and storage networks are Layer 2 fabrics: Ethernet fabrics for server clusters and storage fabrics for shared storage. This may be a good time to consider how the word “fabric” is used in this context. Fabrics have several important properties, including self-aggregation, self-healing, transparent internal topology, and equal-cost multipath routing at Layer 2. For the compute pool, Ethernet fabrics eliminate Spanning Tree Protocol (STP) and manual configuration of Inter-Switch Links (ISLs) and have self-aggregating ISLs (trunks) with load balancing and automatic link failover. They include distributed intelligence so that network policies and devices are known at every port.

When virtual machines move across network ports, port configuration and policies are consistently applied to application traffic. Traffic in Ethernet fabrics is not restricted to “north-south” flows, which move up and down network tiers. Instead, Ethernet networks are flattened so they can efficiently span multiple server racks, reducing latency and avoiding congestion on ISLs.

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 Spanning Tree Protocol, while still being completely interoperable with existing Ethernet networks
- Flexible.** Can be architected in any topology to best meet the needs of any variety of workloads.
- Resilient.** Multiple “least cost” paths are available for high performance and resiliency.
- Elastic.** Easily scales up and down at need.

More advanced Ethernet fabrics borrow further from Fibre Channel fabric constructs:

- 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.
- Protocol convergence (eg Fibre Channel over Ethernet, or FCOE) may also be a feature, intended as a means of better bridging LAN and SAN traffic.

For shared storage, the Fibre Channel storage fabric (SAN) has been an industry standard for more than a decade. With Ethernet fabrics, shared storage is extended to iSCSI and FCoE storage. Virtual machines rely on shared storage networks for their mobility and storage fabrics are the most cost-effective solution whether the storage traffic relies on Fibre Channel or Ethernet for transport.

Client access to applications relies on TCP/IP, HTTP, and HTTPS—which requires Secure Socket Layer (SSL) to secure client connections to applications. An Application Delivery Controller (ADC) pools multiple client connections into one or more virtual machines hosting the virtualized Web, application, and database servers. This balances the workload on physical servers, resulting in improved server efficiency. An ADC also improves server efficiency via SSL offloading, traffic shaping, and load balancing of multiple client connections. In addition, built-in connection and traffic monitoring per application can integrate with virtual machine orchestration software, so virtual machines come online or go offline as client workloads change. To provide cost-efficient scalability, an ADC with fast hardware for scale-up and clustering for scale-out is preferred. Finally, load balancing can extend between data centers directly supporting hybrid clouds so that customer-owned data centers can efficiently use public cloud computing services to handle intermittent spikes in application workloads cost effectively.

With Brocade VCS technology, the cloud-optimized network provides:

- **Unmatched simplicity**
- **Non-stop networking**
- **Optimized applications**
- **Investment protection**

Brocade VCS Technology Provides a Cloud-Optimized Network

With the introduction of Brocade VCS™ fabric technology for the data center, Brocade has transformed classic Ethernet into an Ethernet fabric. As shown in Figure 2 on the following page, a Brocade VCS Ethernet fabric flattens the data center network and connects to the core routers, which can be Brocade MLX Series or other vendor routers. Storage traffic from the Ethernet fabric can access the Brocade DCX® Backbone family of Fibre Channel storage network switches. Via core routers, the Brocade ServerIron® ADX Series of high-performance application delivery controllers provides secure, scalable, and efficient client access to applications. In the future, Brocade VCS Dynamic Services will integrate Brocade ServerIron ADX functionality into the Ethernet fabric along with firewall and encryption services, simplifying placement and management of these essential functions. Brocade VCS Dynamic Services will also provide Ethernet fabric extension between data centers so applications can easily move between private and public cloud services.

With Brocade VCS technology, Brocade is the first vendor to deliver an Ethernet fabric for a cloud-optimized network, a network that is ready for deployment in both public and private cloud data centers.

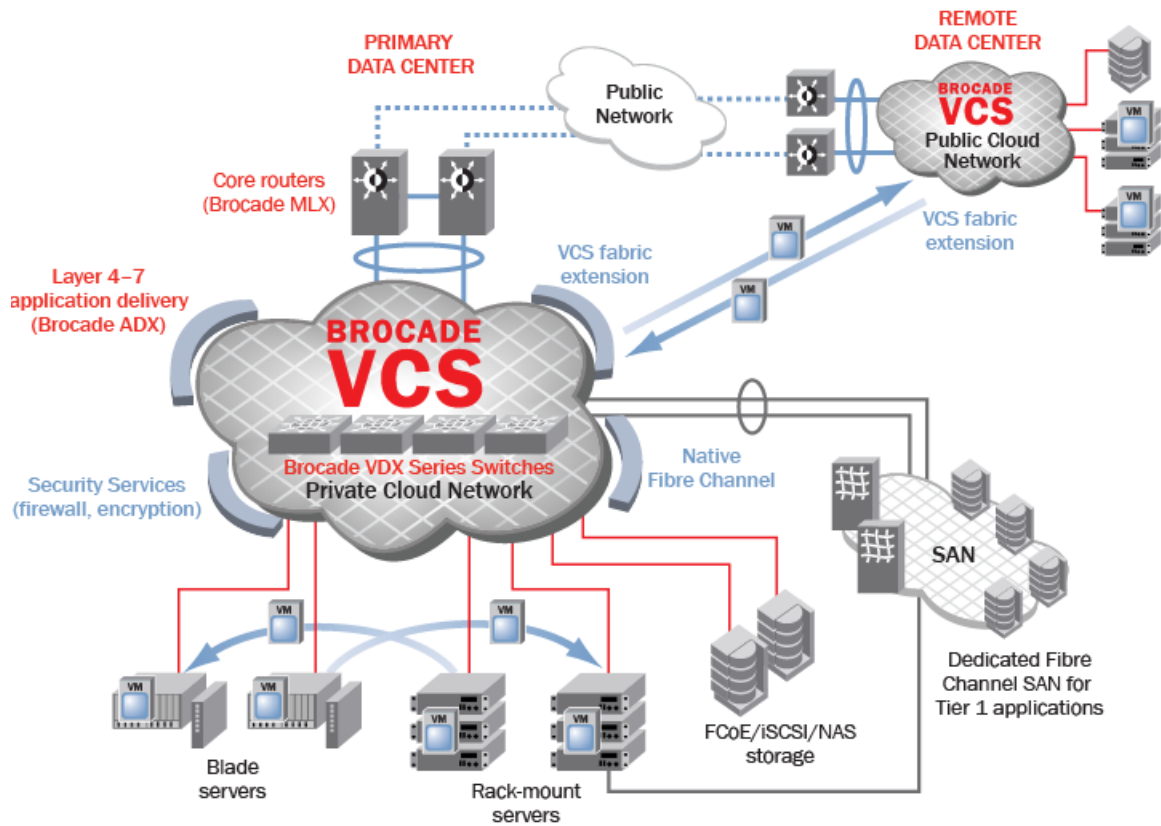


Figure 2. Brocade cloud-optimized network architecture

© 2011 Brocade Communications Systems, Inc. All Rights Reserved. 08/11 GA-PP-350-01-02

Brocade, the B-wing symbol, DCX, Fabric OS, and SAN Health are registered trademarks, and Brocade Assurance, Brocade NET Health, Brocade One, CloudPlex, MLX, VCS, VDX, and When the Mission Is Critical, the Network Is Brocade are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned are or may be trademarks or service marks of their respective owners.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

