

---

# **Microsoft** | Virtualization

## **Simplify Virtual Machine Management and Migration with Ethernet Fabrics in the Datacenter**

*Enabling automatic migration of port profiles under Microsoft Hyper-V with Brocade Virtual Cluster Switching technology*

Solution Blueprint

Published: March 2011

**Abstract:** This white paper provides a blueprint for a virtualized datacenter architecture based on Microsoft® Hyper-V™ and an Ethernet fabric enabled by Brocade Virtual Cluster Switching™ technology. The solution reduces management complexity, provides more efficient use of network resources, and enables easier migration of virtual machines from one physical server to another.

# Copyright

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This White Paper is for informational purposes only. [MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.](#)

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2011 Microsoft Corporation. All rights reserved.

Microsoft, Windows, Windows Server, Hyper-V, and PowerShell are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

# Contents

Contents .....	iii
Executive Summary .....	4
Introduction .....	4
Next-generation virtualization.....	4
Virtual machine migration .....	5
Ethernet network and Spanning Tree Protocol challenges .....	5
Key Concepts.....	7
Private cloud .....	7
Port profiles.....	7
Server virtualization .....	7
Brocade Virtual Cluster Switching technology.....	8
Ethernet fabric.....	8
Distributed Intelligence .....	9
Automatic migration of port profiles.....	9
Logical Chassis .....	9
Solution Components.....	10
Microsoft technologies .....	10
Windows Server 2008 R2 .....	10
Hyper-V Live Migration .....	12
System Center Virtual Machine Manager .....	12
Brocade Technologies .....	13
VDX switches with VCS technology.....	13
Blueprint.....	14
Overview .....	14
Flat network for virtual machine migration between racks use case .....	15
Setup detail.....	15
Conclusion.....	20
Resources .....	20

## Executive Summary

The rapid adoption of virtualization has enabled organizations to consolidate physical servers and make more efficient use of their datacenter investments. Now, many organizations are considering how they can move beyond the first generation of virtualization toward a dynamic datacenter or a “private cloud” model in which computing power is delivered as an elastic, on-demand service with automated management and load balancing. In such a system, applications and operating systems can be moved easily from one server to another depending on changing resource needs. It is possible to implement such a model today using a combination of technologies from Microsoft® and Brocade®.

Microsoft Windows Server® 2008 R2 Hyper-V™ is the hypervisor that supports the live migration of running workloads without user disruption, while Microsoft System Center Virtual Machine Manager (SCVMM) provides unified management of physical and virtual workloads. Brocade Virtual Cluster Switching™ (VCS) technology, available in the Brocade VDX 6720 Data Center Switch family, delivers the first true Ethernet fabric for the datacenter. By eliminating the need for Spanning Tree Protocol (STP), VCS technology enables virtual machines to be moved efficiently within and among racks without the need to manually reconfigure the network or constantly change network policies on switch ports. By flattening the network hierarchy and enabling all switches to be managed as a single logical chassis, VCS technology greatly simplifies the management and optimization of virtualized datacenters running Hyper-V and SCVMM.

## Introduction

Organizations want to exploit the capabilities of virtualization beyond hardware consolidation to increase datacenter automation, scalability and flexibility. Until now, the limitations of traditional Ethernet networks have created challenges to implement these “private cloud” capabilities, particularly when it comes to migrating virtual machines from server to server. Ethernet fabric topology enabled by Brocade VCS technology, working with Microsoft virtualization and management solutions, solve a number of these challenges. One of the key challenges addressed by an Ethernet fabric is the simplification of virtual machine migration between servers and racks.

### Next-generation virtualization

Before virtualization, datacenter workloads and server operating system images were typically provisioned in a one-to-one relationship with physical servers; in other words, one server per workload. While this approach is simple, stable, and secure, it severely underutilizes the power of the server hardware available today. It also hampers agility by requiring organizations to purchase, deploy, and manage additional hardware any time they need to scale or add workloads. With continued rapid growth in computing needs, datacenter operators quickly reach budget limitations, struggle with management complexity, and require more physical space.

These challenges explain the rapid adoption of virtualization, which separates the server operating system and workloads from the underlying hardware. This enables multiple

workloads to be run on a single server. While not all workloads can be virtualized, organizations have realized significant cost and management benefits from implementing hardware virtualization.

Hardware virtualization greatly increases server utilization, and therefore reduces the need to buy and manage additional hardware. However, it is only the first step in realizing the full power of datacenter virtualization. The ultimate vision is a fully dynamic datacenter or “private cloud” that provides the following benefits:

- Servers can host a wide variety of workloads.
- Running workloads can be moved without impacting users.
- Allocated IT resources can contract and expand as needed.
- Moving virtual machines does not require manual network reconfiguration.
- A self-service model enables people to access the resources they need without specialized IT knowledge.

## Virtual machine migration

One of the key enablers of a private cloud infrastructure is the easy migration of virtual machines from one physical server to another. Seamless virtual machine migration provides greater resource management automation in the virtualized datacenter. Live migration<sup>1</sup> enables IT professionals to perform hardware maintenance during business hours without affecting users, and move virtual machines as necessary without extensive downtime planning and preparation. Ultimately, it is best to automate the virtual machine migration so that the monitoring systems can optimize datacenter performance without human intervention.

Virtual machine migration and live migration are enabled at the hypervisor level with Microsoft Windows Server 2008 R2 Hyper-V. Microsoft SCVMM enables unified management of physical servers and virtual machines. However, because certain features of the Ethernet technology are the foundation of datacenter networks, moving virtual machines is not always simple. Ethernet technology was available before virtualization, when there was only one operating system installed per server. It was not designed for a virtualized scenario where workloads move from machine to machine and from switch port to switch port.

## Ethernet network and Spanning Tree Protocol challenges

Classic Ethernet networks depend on the Spanning Tree Protocol (STP). STP works to prevent loops in Ethernet networks, but its hierarchical configuration creates virtual machine migration challenges. In classic Layer 2 Ethernet networks, to create a highly available network, STP designates only one active path between switches through the network while any other path is on standby. While this provides an alternate path, only one path can be used at a time. Because of this, a maximum of only 50 percent of the available network bandwidth can be used. Since one of the goals of server virtualization

---

<sup>1</sup> *Live migration is the process whereby virtual machines can be moved while running without going out of service.*

is to increase physical server utilization, increased network bandwidth utilization should also be expected.

To increase network utilization, Multiple Spanning Tree Protocol (MSTP) and similar protocols enable separate spanning trees per VLAN. While this improves bandwidth utilization, the STP limit of one active path between switches remains. Virtual machines moving between racks must traverse the hierarchical tree of switches instead of going directly to their intended destination. This is very inefficient. In addition, because traffic paths are manually configured with MSTP, the complexity increases.

For a virtual machine to migrate from one server to another, original and destination servers must be configured identically in many respects. This principle applies to the network as well. VLAN, Access Control List (ACL), Quality of Service (QoS), and so on must be identical on both the source and destination access switch ports. If switch port configurations differ, either the pre-migration testing will fail or network access for the virtual machine will break. Organizations could map all settings to all network ports, but this would violate most networking and security best practices.

With STP, any change in the physical topology of links requires a new spanning tree. As a new spanning tree is established, all traffic on the LAN stops. This can take anywhere from five seconds with Rapid Spanning Tree Protocol (RSTP) or up to several minutes with STP. Please note that this convergence can vary unpredictably even with small topology changes. The demands for non-stop traffic flow increases with server virtualization, and consequently network convergence times have to shrink. STP does not provide an adequate solution for these requirements. Finally, when a spanning tree is reestablished, broadcast storms can occur and cause a network slowdown. Layer 2 networks are typically minimized in the datacenter due to these STP limitations and therefore, virtual machine scalability and mobility is limited.

To address these issues and enable the advancement of the dynamic datacenter, Brocade introduces Brocade Virtual Cluster Switching (VCS), a new Layer 2 Ethernet technology that improves network utilization, maximizes application availability, increases scalability, and dramatically simplifies the network architecture in any datacenter. VCS technology provides the first true Ethernet fabric which effectively addresses the challenges of virtualized datacenters.

VCS technology makes virtual machine migration much easier because it:

- Flattens the network for more efficient east-west traffic flow.
- Makes all VCS-enabled switches aware of virtual machine locations and consistently applies network policies.
- Does not require manual intervention when a virtual machine moves.
- Automatically uses least-cost paths for low latency, maximum bandwidth utilization, and path resiliency.
- Supports heterogeneous server virtualization in the same network.

## Key Concepts

This section reviews concepts and technologies that underpin the solution blueprint, including private clouds, port profiles, server virtualization, and the essential features of Brocade VCS technology.

### Private cloud

The basis of this blueprint combines virtualization with Ethernet fabrics to enable a datacenter with many of the characteristics of a private cloud. The characteristics of a private cloud include:

- **Scalability:** Computing power and storage can meet increased demand on a just-in-time basis.
- **Elasticity:** Computing resources expand and contract automatically to precisely meet needs at a given point in time.
- **Multi-tenancy:** Consistent policies per virtual machine ensure secure multi-tenancy.
- **Self-service:** Users can access the resources they need without intervention from IT administrators.

One of the key ways that the solution achieves these goals is by removing the need to manually reconfigure the network when migrating virtual machines.

### Port profiles

In a classic Ethernet, network policies and security control are assigned at the physical switch port. They are applied to all traffic entering the switch port and are tied to the physical port. Traffic is identified by the MAC address of the sending device, so a network policy includes a MAC address.

A port profile separates network policies from the physical port, acting as a portable container for a variety of policies. With VCS technology, a port profile can include one or more of the following profiles: VLAN, QoS, FCoE, and security. In a VCS Ethernet fabric, all switches and ports access a common port profile database so they automatically and consistently apply the correct set of port profiles to a virtual machine MAC address even when that virtual machine moves to another port on a different switch in the Ethernet fabric. Network management issues limiting virtual machine mobility are eliminated.

### Server virtualization

By enabling organizations to run multiple operating systems and applications on a single physical server versus multiple physical machines, server virtualization can help organizations reduce hardware, energy, and management overhead significantly. Server virtualization is also a key building block of a dynamic IT infrastructure which is all about efficiency, resiliency, and agility.

From an application perspective, server virtualization enables organizations to consolidate workloads. This increases resource utilization while reducing capital and

operational costs. In addition, server virtualization can help organizations improve application performance, availability, management, and agility. The more agile the IT environment is, the more an organization is able to meet changing business and application requirements.

## Brocade Virtual Cluster Switching technology

The key technology pillars of VCS technology are Ethernet fabric, Distributed Intelligence, and Logical Chassis. The following information describes how each pillar supports enhanced virtual machine migration as well as better performance and manageability of virtualized datacenter environments.

### Ethernet fabric

The Ethernet fabric provided with VCS technology includes advanced functions such as the emerging standard called Transparent Interconnect of Lots of Links (TRILL)<sup>2</sup>. Like STP, TRILL prevents loop formation. With TRILL, however, all paths in the network are active and traffic is distributed across those equal cost paths automatically. In this optimized environment, traffic automatically takes the shortest path for minimum latency without manual configuration requirements. Additionally, because it follows the standards for Data Center Bridging (DCB), the fabric is lossless with low latency, so it makes efficient use of network resources and can support storage traffic such as Fibre Channel over Ethernet (FCoE) that requires lossless Ethernet for transport. This becomes critical as network utilization increases with the many-fold increase in workloads per server, common to virtualized datacenters.

The Ethernet fabric is self-forming. When two VCS-enabled switches are connected, an inter-switch link (ISL) is automatically created and the switches learn the common fabric topology. The Ethernet fabric does not dictate any specific topology, so it does not restrict over-subscription ratios. This allows the architect to create a topology that best meets application requirements—especially important with the additional architectural complexity of virtualized environments. Before virtualization, servers carried a predictable amount of traffic. The over-subscription ratio could be designed into the network topology with a relatively low risk of unexpected congestion. In any case, once the ratio was set, it could be difficult to change. To change the ratio required reconfiguring links between access switches and Link Aggregation Groups (LAGs), between the access layer and aggregation layer, and between the aggregation and core layers so reconfiguration cascaded across the network.

In a virtualized datacenter, all bets are off. Bandwidth needs to be able to scale quickly to meet the changes in application requirements. Traffic can shift dramatically and therefore, lower oversubscription ratios are required. The VCS Ethernet fabric includes Brocade ISL Trunks which aggregate multiple ISLs into a single resilient logical

---

<sup>2</sup> VCS technology currently uses a pre-standard implementation. Brocade will extend VCS to include more TRILL and associated standard functionality as it is approved by the various standards bodies. Enhancements will be available through firmware upgrades.

connection. This enables the oversubscription ratio to be set by adding another link—in other words, simply plugging in a cable—so the Brocade ISL Trunk has more bandwidth.

The Ethernet fabric is much more robust and resilient than Ethernet networks using STP. Events like added, removed, or failed links are not disruptive to the traffic moving across the Ethernet fabric and do not require all traffic in the fabric to stop. Inter-switch links automatically configure and also automatically form Brocade ISL Trunks for increased bandwidth, high performance load balancing in hardware and link resiliency. If a single link fails, traffic is automatically rerouted to other available paths quickly enough that even sensitive storage traffic is unaffected. Switch or port failures require the fabric to learn the new topology, but this does not halt existing traffic flows ensuring that unaffected traffic continues to flow. This allows inter-switch bandwidth to be quickly changed to meet new needs simply by adding or removing switches and/or cables.

In a dynamic virtualized environment, an Ethernet fabric makes it easy to move virtual machines, not only between servers in a single rack, but from rack to rack. Where STP would require the migration traffic to traverse a hierarchical tree to move between racks, with an Ethernet fabric, the traffic can go directly to its destination using the shortest possible path, and with no manual reconfiguration required.

### **Distributed Intelligence**

With VCS technology, all configuration and end device information is automatically distributed to each switch in the fabric. The fabric is aware of all switches and connected devices including virtual machines. When a virtual machine connects to the fabric for the first time, all switches in the fabric learn about that virtual machine. This enables fabric switches to be added or removed, and for physical or virtual servers to be relocated without manual reconfiguration of the fabric.

### **Automatic migration of port profiles**

Most importantly for the purposes of this blueprint, the Automatic Migration of Port Profiles (AMPP) service that is part of VCS Distributed Intelligence supports virtual machine migrations to another physical server, ensuring that the ingress ports will have the same policies applied to the traffic from the virtual machine. If a virtual machine's MAC address moves from one switch port to another, the switch has access to a common database of port profiles which it can use to determine the correct policies to apply to the traffic. The Virtual Machine Manager can automatically move a workload based on the needs of the system as a whole without concern for network reconfiguration.

### **Logical Chassis<sup>3</sup>**

All switches in a VCS Ethernet fabric are managed as if they were in a single logical chassis. To the rest of the network, the entire fabric looks no different than any other single layer 2 switch, whether it contains a few ports or thousands. The VCS Ethernet

---

<sup>3</sup> Roadmap item is not available in the initial release.

fabric is designed to scale without increasing management complexity. Consequently, VCS technology removes the need for separate aggregation switches because the fabric is self-aggregating. This enables the network architecture to be flattened, dramatically reducing cost and management complexity.

With the VCS Logical Chassis, each physical switch in the fabric is managed as if it were a port module in a chassis. This enables fabric scalability without manual configuration. When you add a port module to a chassis, just plug it into the chassis, and a switch can be added to the Ethernet fabric just as easily. The logical chassis functionality drastically reduces management of small form-factor Top of Row (ToR) switches. The ToR switches are not managed individually, but rather as a single logical chassis.

## Solution Components

This section discusses the specific software and hardware products from Microsoft and Brocade that make up the solution.

### Microsoft technologies

Microsoft virtualization provides an architecture that can profoundly affect nearly every aspect of the IT infrastructure management lifecycle. It can drive greater efficiencies, flexibility, and cost-effectiveness throughout the organization.

A standard Microsoft virtualization implementation is typically structured using the Windows Server 2008 R2 Hyper-V role to enable virtualization and Windows Server Failover Cluster to handle high availability and disaster recovery requirements. System Center Virtual Machine Manager (SCVMM) is typically used to simplify virtualization management.

### Windows Server 2008 R2

Windows Server 2008 R2 builds on the foundation of Windows Server 2008, expanding existing technology and adding new features to enable organizations to increase the reliability and flexibility of their server infrastructures. Powerful tools such as Internet Information Services (IIS) version 7.5, updated Server Manager and Hyper-V platforms, and Windows PowerShell™ version 2.0 combine to give customers greater control, increased efficiency, and the ability to react to front-line business needs faster than ever before.

<b>Features</b>	<b>Benefits</b>
Built-in Windows Server 2008 R2 Hyper-V technology	Enables enterprises to easily leverage the benefits of virtualization without adopting a new technology.
Broad device driver support	The new 64-bit micro-kernelized hypervisor architecture leverages the broad device driver support in the Windows Server 2008 R2 parent partition to extend support to a broad array of servers, storage types, and devices.
Symmetric Multiprocessor (SMP) support	Hyper-V supports SMP on virtual machines.
Host high availability	Windows Server 2008 R2 clustering provides high availability to virtual machines to minimize unplanned downtime.
Shared storage high availability	Microsoft MPIO dynamically routes I/O to the best path and protects against connection failures at any point between a Hyper-V host and shared storage including NICs/adapters, switches, or array ports.
Live virtual machine migration	Live migration is available to support business continuity during planned downtime and over distances.
Volume Shadow Copy support (VSS)	Provides a robust host-based API to enable virtual machine backup by leveraging the existing Windows VSS-based infrastructure.
Easy extensibility	Easy extensibility is available by using the standards-based Windows Management Instrumentation (WMI) interfaces and PowerShell.
Simplified integrated management	With its tight integration into the Microsoft System Center family of products, customers have end-to-end physical and virtual infrastructure management capability for Hyper-V environments.

## Hyper-V Live Migration

Hyper-V Live Migration is a key enabling technology for the private cloud as envisioned by this blueprint. It supports the movement of running virtual machines from one Hyper-V physical host to another without any disruption of service or perceived downtime. Datacenters with multiple Hyper-V physical hosts will be able to move running virtual machines to the best physical computer for performance, scaling, or optimal consolidation without impacting users. Live migration makes it possible to keep virtual machines online, even during maintenance, increasing productivity for both users and server administrators. Datacenters will be also able to reduce power consumption by dynamically increasing consolidation ratios and powering off idle physical hosts during times of lower demand.

## System Center Virtual Machine Manager

For IT professionals responsible for managing virtual infrastructures, System Center Virtual Machine Manager (SCVMM) 2008 R2 provides a cost-effective solution for the unified management of physical machines and virtual machines.

Features/Capabilities	Benefits
Enterprise-class management suite	Manages both Hyper-V and VMware® ESX® virtualization environments.
Intelligent virtual machine placement	Provides support for the intelligent placement of virtual machines.
System Center Operations Manager 2007 R2 integration	Integrates with System Center Operations Manager 2007 to provide proactive management of both virtual and physical environments through a single console.
Performance and Resource Optimization (PRO)	Ties specific alerts from System Center Operations Manager 2007, such as when hardware fails or load thresholds are exceeded, to remediation actions in SCVMM.
Native P2V/V2V migration	Provides native capability for physical-to-virtual migrations and virtual-to-virtual migrations.
Failover integration	Provides integration with failover clustering to support high availability and the live migration of virtual machines.
Automation	Provides easy automation capabilities leveraging Windows PowerShell.

Features/Capabilities	Benefits
Microsoft System Center Virtual Machine Manager Self Service Portal 2.0	The self-service portal, built on top of Windows Server Hyper-V and System Center, is a free, extensible, turnkey solution that empowers datacenters to dynamically pool, allocate, and manage resources to enable private cloud computing on premises.

### Brocade Technologies

Leveraging Brocade VCS™ technology, Brocade VDX Data Center Switches provide the foundation for the Ethernet fabric—revolutionizing the design of Layer 2 networks and enabling cloud-optimized networking.

#### VDX switches with VCS technology

Brocade VDX Data Center Switches are specifically designed to improve network utilization, maximize application availability, increase scalability, and dramatically simplify network architecture in virtualized datacenters.

Features	Benefits
Ethernet fabric	Enables the construction of a flat, multi-path, deterministic fabric for the datacenter that is lossless, low latency, multi-pathing, and able to forward any type of traffic. Traffic automatically travels along the lowest cost path, lowering latency and efficiently using all available bandwidth.
Distributed Intelligence	Each VCS-enabled switch shares information about every attached device. Device policies are managed from a common database or “port profile” that allows policies to be dynamic so they move with the device MAC address. Distributed Intelligence also allows the fabric to be self-forming, enabling bandwidth to be added simply by adding another link between switches.
Logical Chassis	VCS dramatically reduces management complexity and cost by enabling the entire fabric to be treated as one logical switch.

Data Center Bridging (DCB)	Supports the required extensions to Ethernet making it lossless. This is essential when block storage traffic such as Fibre Channel over Ethernet (FCoE) and iSCSI traffic flow across the VCS Ethernet fabric.
Built for flexibility	Each port supports 1 GbE and 10 GbE as well as a variety of connection types including direct-attached copper, direct-attached optical, and fiber-optic cabling.
Dynamic Services	In the future, services such as fabric extension over distance, application delivery control and native fibre channel connectivity can be added. Through VCS, the new switches and software with these services behave as service modules within a logical chassis. Furthermore, the new services are then made available to flow within the fabric, dynamically evolving the fabric with new functionality.

## Blueprint

The blueprint provides details on how to set up the VDX 6720 switches and a VCS Ethernet Fabric using the equipment and software described in previous sections. This configuration enables easier migration of virtual machines between racks.

### Overview

The proof of concept consists of four server racks with an Ethernet fabric of 10 VDX 6720-60 switches shown in Figure 1. The operating system is Windows Server 2008 R2 with Hyper-V employed as the hypervisor. The system is managed using Microsoft SCVMM. Brocade VDX switches use an industry-standard command line interface while the SCVMM graphical user interface is used to manage the server and the Hyper-V virtual machines.

This environment provides the key functions of a network in which the movement of virtual machines is frictionless, requiring no manual reconfiguration of the network. The network must be configured so there is sufficient ISL bandwidth available for workloads to move across switch ports. The topology must be designed to minimize congestion and utilize a cut-through, non-blocking switch architecture rather than the store-and-forward architecture that is common to low-end Ethernet devices. The network must provide linear scaling as servers are added to racks and virtual machines are added to servers.

## Flat network for virtual machine migration between racks use case

In this blueprint, we include a use case that flattens the network and eliminates network policy configuration problems when virtual machines migrate. For brevity, only the required network configuration commands are listed. For more detail about these commands, refer to the Brocade Network Operating System (NOS) Command Reference Guide available at [my.brocade.com](http://my.brocade.com).

In this use case, the eight total top-of-rack switches and two middle-of-row switches are managed as a single VCS Logical Chassis as shown in Figure 1. This enables seamless movement of virtual machines among the servers in any of the four racks.

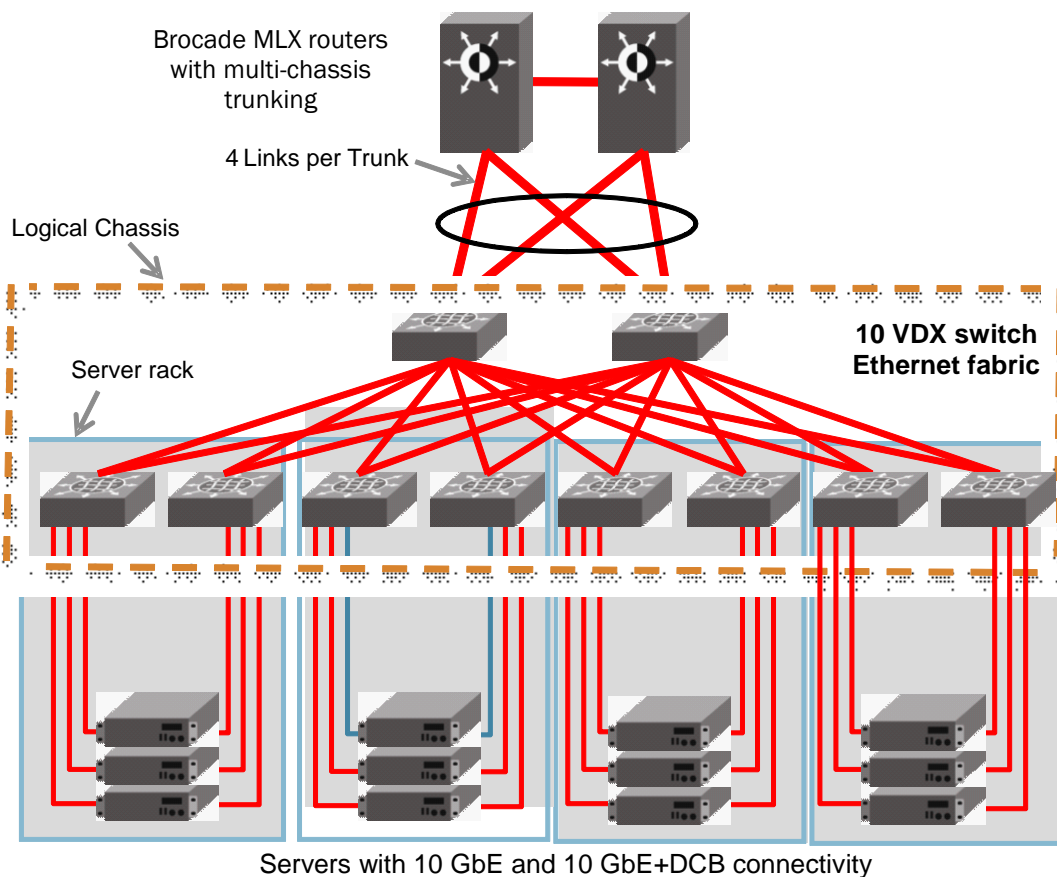


Figure 1: Flat network for virtual machine migration between racks

### Setup detail

The following steps illustrate the process of setting up an Ethernet fabric.

#### 1. Select appropriate network architecture.

Select a fabric architecture that can accommodate growth appropriately. A good design principle is to leave flexibility in the solution to accommodate planned and unplanned

growth without needing to change the topology or the fabric architecture. The topology selected is shown in Figure 2. It consists of a 10-node VCS fabric using 10 60-port VDX switches. They are arranged in a Clos architecture, with redundant VDX switches at the middle-of-row functioning as the aggregation layer.




Figure 2: Scalable fabric topology

## 2. Configure and implement VCS.

ISLs between the VDX switches auto-configure to form an Ethernet fabric. Adjacent ports within a port group automatically create a Brocade ISL Trunk inside the VCS Ethernet Fabric. Hardware-based frame forwarding provides high link utilization across all ISL links as well as automatic link resiliency.

As shown in Figure 3, each rack can support up to 36 servers connecting to each ToR VDX 6720 switch. The ToR switch then connects to each of the aggregation VDXs with four ISLs per Brocade ISL Trunk. Because data takes the shortest, least-cost paths and paths can forward traffic, this topology provides 40 Gbps of bandwidth per trunk, and a total of 160 Gbps per rack to the core.




Figure 3: Topology for implementing VCS technology


### 3. Configure host port access.

When connecting each server, the access port must be configured in accordance with the VLAN topology defined for the datacenter and application infrastructure. In this use case, each host port is placed in VLAN 100 which is configured on the host-port interface. The following configuration commands are used:

- Configure terminal
- Interface TenGigabitEthernet 1/0/5
- Switchport
- Switchport mode access
- Switchport access vlan 100

### 4. Determine virtual machine policies and create profiles.

Based on the service level agreement for each virtual machine, policies are defined and port profiles are created and applied. As shown in Figure 4, the VCS port profile is a container where policies are organized by four classes of profiles including FCoE, VLAN, QoS, and security. A port profile container can have one or more profiles as required. A port profile is assigned to one or more device MAC addresses so that all the profiles within the port profile container are applied to the traffic originating from them.



**Figure 4: Profile classes inside the port profile container**

In this example, a security network profile is applied for Tier X applications using the following steps:

- a. Using the following commands, configure the Access Control List (ACL) to be used in the port profile. This is called "TierX\_App\_acl". In this example, traffic to MAC "0050.563f.0001 0050.563f.9999" is not allowed.

```
configure terminal
mac access-list extended TierX_App_acl
seq 10 deny 0050.563f.0001 0050.563f.9999 count
seq 20 permit any any
```

- b. Using the following commands, create the port profile "TierX\_App" including the acl.

```
port-profile TierX_App
security-profile
mac access-group TierX_App_acl in
vlan-profile
switchport
switchport mode access
switchport access vlan 100
```

- c. Display the complete profile to confirm your entries. In this case, use the following commands to show the port-profile named TierX\_App.

```
port-profile TierX_App
ppid 1
vlan-profile
switchport
switchport mode access
switchport access vlan 100
security-profile
mac access-group TierX_App_acl in
```

- d. Multiple port profiles are stored in a common database. Therefore, it is necessary to activate a profile before using it as shown in the following command:

```
port-profile TierX_App activate
```

- e. Next, associate the port profile with the MAC addresses of the appropriate virtual machines using the following commands:

```
port-profile TierX_App static 0050.563f.0001
port-profile TierX_App static 0050.563f.0002
port-profile TierX_App static 0050.563f.0003
port-profile TierX_App static 0050.563f.0004
```

- f. The following commands enable you to display a listing of the port profile and current status to confirm the correct configuration:

```
show port-profile
show port-profile status
```

- g. Finally, for the edge ports requiring a port profile, indicate that ingress traffic should be filtered by port profiles. You can turn on Port Profile filtering on some or all edge ports as an option with the following commands:

```
int te 1/0/6
no switchport
port-profile-port
int te 2/0/6
no switchport
port-profile-port
```

Now that the port profile is configured, associated with virtual machine MAC addresses and the appropriate edge ports are filtering traffic using port profiles, Automated Migration of Port Profiles (AMPP) is operational.

The port profile container and its associated polices are consistently enforced as a virtual machine moves across the edge ports of the VCS Ethernet Fabric and are agnostic to the hypervisor technology in use. Figure 5 shows a virtual machine moving from one clustered server to another. Traffic was prohibited to a particular MAC as shown by the red X and continues to be blocked when the virtual machine moves to the destination server.




Figure 5: Virtual machine moving from one clustered server to another.

## Conclusion

With increasing demands being put on today's datacenters, server consolidation is only the first step in using virtualization to its fullest. By combining the built-in virtualization capabilities of Windows Server 2008 R2, the centralized management capabilities of SCVMM, and the revolutionary advancements of Brocade VCS technology, organizations can take a step closer to the efficiency, cost, and scalability benefits of the private cloud. This blueprint has demonstrated how an Ethernet fabric-enabled network can enable faster, easier, more automated migration and management of virtual machines both within and between server racks. And this only represents a small part of what these technologies can do together to enable the next step in datacenter evolution.

## Resources

For more information on the technology components of the solution detailed in this blueprint, see the following:

Brocade Solutions for Microsoft: <http://www.brocade.com/microsoft>

Microsoft Virtualization Solutions: <http://www.microsoft.com/virtualization>

Brocade VDX 6720 Data Center Switches with VCS Technology:  
<http://www.brocade.com/vdx>