NORTEL

Solution Brief
Resilient data center server edge solutions for VMware Virtual Infrastructure 3

Market requirements for a virtual server network

As customers realize the benefits of machine virtualization in the data center, it is clear that optimizing resource usage and minimizing server and network downtime are key requirements. As an individual physical server can have many virtual machines running critical applications, both the server hardware and the networking topology must eliminate any single points of failure, provide instant failover capability for link or node failures, and ideally maximize the use of deployed switches, network interface cards (NICs) and cabling. Furthermore, it is imperative that both planned and unplanned downtime be minimized — specifically, the data center operations staff should be able to perform any maintenance tasks on the network and servers without impacting service.

VMware Virtual Infrastructure Solution

VMware is a clear leader in virtual server solutions in the marketplace today. VMware Infrastructure 3 (VI3) offers the ability to create multiple virtual machines on the same server to better utilize server platforms. In addition to server consolidation, VMware VI3 offers a complete virtualized infrastructure solution for business continuity encompassing data protection, high availability and disaster recovery. This solution translates into reduced hardware and operating costs, reducing the time it takes to provision new servers and saving more money per year for every server workload virtualized.

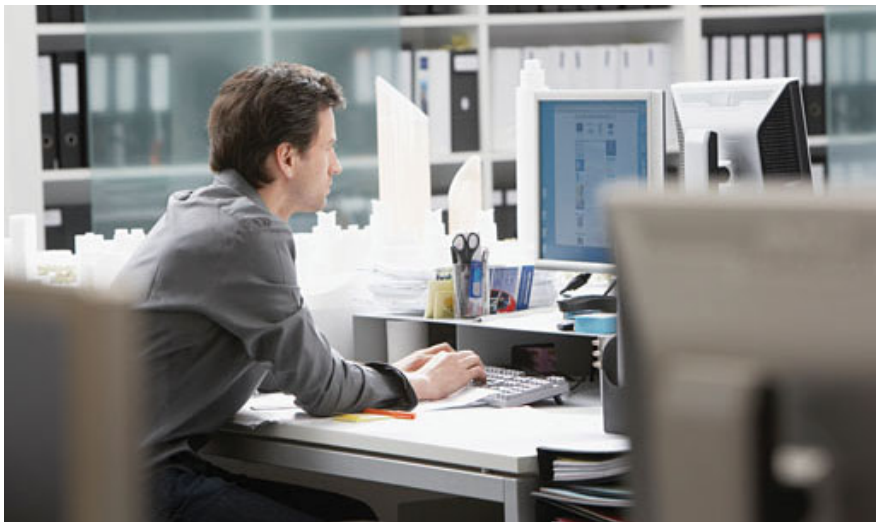

VMware VI3 incorporates VMware ESX Server 3 platform as the foundation for hosting guest operating systems. VMware virtual networking recommendations include:

- Using 802.1Q VLAN trunks and Virtual Switch Tagging (VST) for VLAN assignment and isolation
- Logical separation via VLANs of Service Console, VMotion, iSCSI and Virtual Machine traffic
- Using NIC teaming across on board NICs and PCI NICs for resilience and load balancing

In an environment such as iSCSI where multiple iSCSI targets are used, VMware recommends using ESX NIC teams in "IP Hashing" mode to balance traffic across NIC teams and take advantage of multiple switch redundancy.

vmware® | technology alliance PARTNER

Edge switch technology solution

There are several options to connect an edge switch to a VMware ESX server to achieve high-density Gigabit Ethernet access. Whether you're connecting a single switch, two switches for redundancy or a stack of switches for connecting hundreds of ESX servers, the solution must offer simplicity, resiliency, scalability, performance and the best total cost of ownership (TCO).

Currently, other solutions require network designers to either choose between resiliency and performance or to add cost and complexity. With active-standby style resiliency, failover time is increased and network equipment, cabling and server NICs are underutilized. Without a high-performance stacking solution, leveraging uplinks across the stack is not practical. And most importantly, without switch clustering, network maintenance and operations can impose a costly service outage.

Introducing Nortel's solution

Nortel offers unique technologies to meet the simplicity, resiliency, scalability and performance while providing lowest TCO.

Switch clustering

Split Multi-Link Trunking (SMLT) is a Nortel architecture that helps eliminate single points of failure and creates multiple paths from user access switches to the core of the network. Compatible with 802.3ad, SMLT does more than prevent network loops. SMLT provides an architecture to design resiliency directly into the network. It also works to reroute failures as quickly as possible. In most cases, network reconvergence is sub-second. Nortel's SMLT is an extension to the IEEE 802.3ad link aggregation specification. SMLT avoids loops due to its superior enhanced link aggregation control protocol. If 802.1d is used, multiple Spanning Tree groups are required and VLANs must be manually assigned to those groups — all of which makes ongoing administration and troubleshooting extremely complex. With SMLT, it is no longer necessary to use the Spanning Tree protocols to design resilient networks. SMLT provides much faster convergence times than Spanning Tree (typically one second versus 30 to 60 seconds). SMLT also eliminates the blocking of ports by Spanning Tree protocols, thus increasing network bandwidth since all links in a trunk can be utilized for forwarding traffic.

SMLT allows two aggregation switches to appear as a single device to dual-homed switches and servers. The aggregation switches make use of an InterSwitch Trunk (IST) over which they exchange information, permitting rapid fault detection and forwarding path modification. To achieve network element protection, SMLT extends link aggregation to allow dual homing of IEEE 802.3ad attached devices. Both of the dual-homed connected devices are active and pass traffic. This architecture provides twice the available bandwidth of merely using the Spanning Tree Protocol.

SMLT improves the reliability of a Layer 2 network operating between the user access switches in a building and the network center aggregation switch, as well as with the connections to multi-homed servers. It does so by providing load sharing among all available links and fast failover in the case of a link or core switch failure.

Nortel Switch Clustering (SMLT) is a solution provided on a number of Ethernet Routing Switches including modular switches Ethernet Routing Switch 8600 and Ethernet Routing Switch 8300 as well as fixed switches Ethernet Routing Switch 1600 and Ethernet Routing Switch 5500.

Horizontal stacking

Horizontal stacking uses a unique application of the stack cables to connect multiple top-of-row Ethernet Routing Switches and extend the stack across a number of server racks or cabinets as shown in Figure 1.

Applications such as high-performance computing require connecting hundreds of servers at a low latency with fault-tolerant capability. Of course, the challenge is to do it in a cost-effective way. This means spreading data across hundreds of servers through top-of-row switches without utilizing 10G uplinks

Figure 1. Scalable multi-link trunking into switch

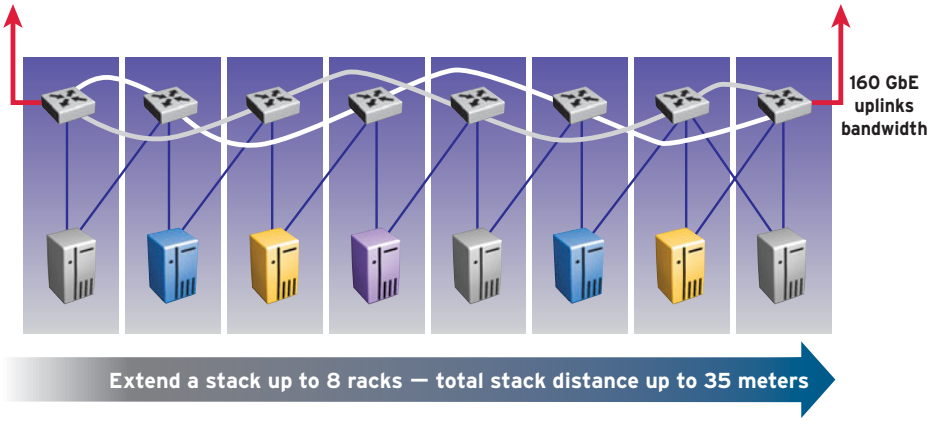

Figure 2. Combining switch clustering and horizontal stacking

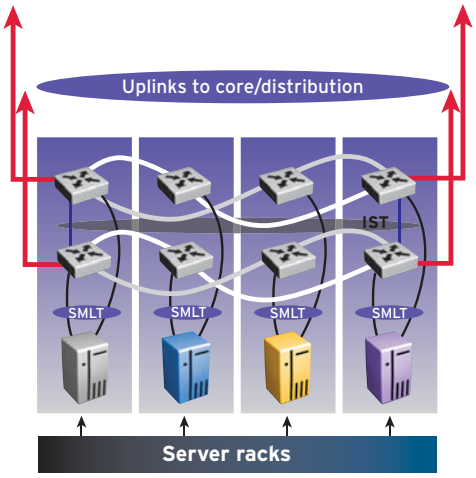

in distribution or core switches. This can only be achieved if stacking bandwidth is high. Nortel's Ethernet Routing Switch 5500 Series can deliver up to 640-Gbps stacking bandwidth per eight-unit stack. This is ideal for VMotion and iSCSI as it provides very high node-to-node communications bandwidth without impacting uplink capacity.

Key features of horizontal stacking:

- Low latency between servers (9µs)
- Up to 80-Gbps stacking bandwidth per switch
- Up to eight units per stack
- Highly resilient stacking technology with scalable uplinks

Combining switch clustering and horizontal stacking

Combining switch clustering and horizontal stacking gives the ability to connect large numbers of servers with dense in-rack Gigabit connections with active/active links and all of the resiliency benefits of clustering. This architecture offers sub-second failover when it comes to software upgrades, link failure or unit failure. Your edge network is fully protected, with no single point of failure and always on for mission-critical applications. In addition, ongoing maintenance for updating unit configuration or unit replacement can be done without disruption to network operations.

Key advantages of combining switch clustering and horizontal stacking:

- Up to 380 10/100/1000 dual-homed servers
- 640Gbps per stack bandwidth
- Up to 32 x 10Gbps of uplink capacity
- One to eight server racks
- Reduced planned down-time for software upgrades
- Reduced number of uplinks to distribution or core
- Almost zero down-time; bullet-proof architecture
- Simplified maintenance, configuration and updates

Choosing the right switch technology solution – Ethernet Routing Switch 5500

Nortel recommendations for connecting VMware ESX include:

- Use Ethernet Routing Switch Clustering (SMLT) to the server NIC level
- Use horizontal stacking and switch clustering to simplify cabling and provide zero service impact maintenance
- Use ESX NIC teams in IP Hashing mode to balance traffic across NIC teams and take advantage of multiple switch redundancy

The Ethernet Routing Switch 5500 can deliver on these recommendations to provide a low-cost, high-capacity, feature-rich data center edge solution without compromising on resiliency.

Ethernet Routing Switch 5500 delivers a full eight-unit stack with 384 ports that could support 380 servers with the second NIC connected from each server to the other eight-unit stack. 190 servers can be connected in four NIC configurations with a dual eight-unit stack.

Summary

When you bring together industry-leading stacking, performance and active/active resiliency with the simplification provided by Nortel stacking and clustering, you have the absolute best solution for connecting VMware ESX servers in the data center.

Figure 3. Switch technology solution with ERS 5500

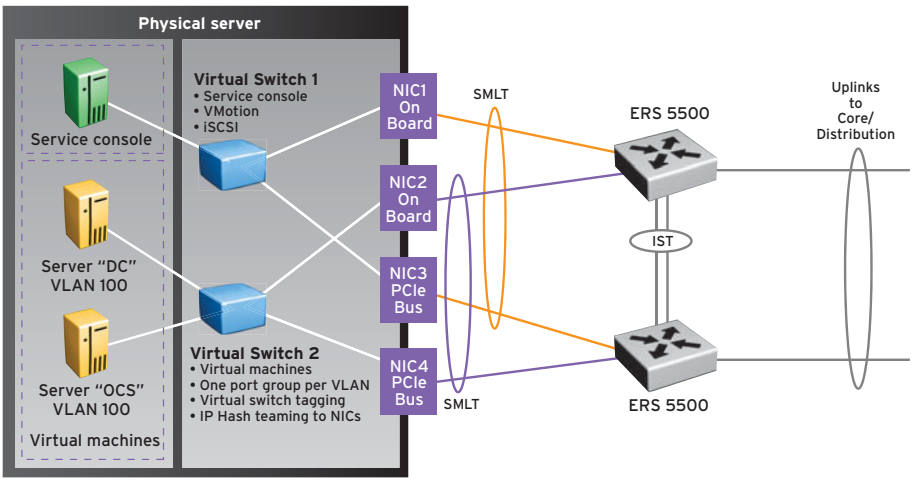

In the United States:
Nortel
35 Davis Drive
Research Triangle Park, NC 27709 USA

In Canada:
Nortel
195 The West Mall
Toronto, Ontario M9C 5K1 Canada

In Caribbean and Latin America:
Nortel
1500 Concorde Terrace
Sunrise, FL 33323 USA

In Europe:
Nortel
Maidenhead Office Park, Westacott Way
Maidenhead Berkshire SL6 3QH, UK
Phone: 00 800 8008 9009

In Asia:
Nortel
United Square
101 Thomson Road
Singapore 307591
Phone: (65) 6287 2877

Nortel is a recognized leader in delivering communications capabilities that make the promise of Business Made Simple a reality for our customers. Our next-generation technologies, for both service provider and enterprise networks, support multimedia and business-critical applications. Nortel's technologies are designed to help eliminate today's barriers to efficiency, speed and performance by simplifying networks and connecting people to the information they need, when they need it. Nortel does business in more than 150 countries around the world. For more information, visit Nortel on the Web at **www.nortel.com**. For the latest Nortel news, visit **www.nortel.com/news**.

For more information, contact your Nortel representative, or call 1-800-4 NORTEL or 1-800-466-7835 from anywhere in North America.

NN123441-022508

NORTEL

BUSINESS MADE SIMPLE