Gartner Research

Publication Date: 7 August 2009 ID Number: G00167903

Use Top-of-Rack Switching for I/O Virtualization and Convergence; the 80/20 Benefits Rule Applies

Joe Skorupa, Mark Fabbi, Robert E. Passmore

This document evaluates the benefits of server input/output (I/O) virtualization and convergence and converged data center networks. It provides a road map for adoption of these approaches. These technologies are immature and care must be taken to avoid proprietary or sub-optimal approaches that will decrease business agility and drive up the total cost of ownership.

Key Findings

- Top-of-rack convergence provides the largest benefit with the least disruption.
- Using a single technology for LAN and storage area network (SAN) access within the rack can deliver savings without requiring a converged backbone.
- Existing solutions may lack sufficient throughput for SAN traffic, forcing Balkanization of storage and driving up cost and complexity.
- Although the T11 Fibre Channel over Ethernet (FCoE) standard has been approved, it is only a piece of the standards puzzle required to construct converged networks.
- Alternative technologies may offer superior approaches to FCoE for in-the-rack networking.

Recommendations

- Consider all of the options for I/O virtualization.
- If you decide to adopt FCoE/Converged Enhanced Ethernet (CEE) as your I/O virtualization technology, use a phased approach to deployment to avoid vendor lock-in and increased long-term costs.
- Require all vendors to contractually commit to no-cost upgrades to the appropriate standards for all components within six months of the ratification of those standards.
- Be wary of vendor claims that you can build a standards-compliant end-to-end converged data center network today, or within the next 18 months.

The Ethernet Steamroller

Once again the industry is abuzz with the promise of a single converged network infrastructure, this time in the data center core. Variously described as FCoE, Data Center Ethernet (DCE), or, more precisely, CEE, this latest set of developments promises to unify computing, networking and storage. This promise is worthy of Tolkien:

"One Infrastructure to rule them all, One Infrastructure to find them, One Infrastructure to bring them all and in the data center bind them!"

While the architectural purity and promise of a single unified network is appealing, Gartner's research demonstrates that such a broad-reaching change is inadvisable at this time and may never be advisable. That said, is there some value to be gained from the work on CEE? Gartner believes that at least 80% of the theoretical value of a completely converged data center network can be achieved over the next 12 to 18 months by limiting the scope of the efforts to the connectivity that is resident within the rack.

The Drivers for Convergence

I/O Consolidation and Virtualization

The hardware requirements for server virtualization are driving the current server upgrade cycle. Additionally, many data centers are moving to high-density blade systems in an effort to reduce the space required to house their continually expanding server requirements. This increases the bandwidth required in each rack and at each switch on the network. Gartner refers to this phenomenon as the increase in "bandwidth density."

As more operating systems/applications are deployed on these new multicore platforms the I/O requirements grow dramatically. In many cases, a 1 rack unit (RU) or 2RU server may require as many as six or eight network ports, or even more, to support the following requirements:

- Two ports of 4 Gbps Fibre Channel (FC) for SAN attachment.
- Two ports of 1 Gbps Ethernet for LAN attachment.
- One port of 1 Gbps Ethernet for VMotion.
- One port of 1 Gbps Ethernet for out-of-band management.
- One port of 1 Gbps Ethernet for backup.

This leads to significant power and cooling requirements just to support server adapters and their associated switch ports (as much as 10% of total server power, or more). Additionally, cable management for a rack containing 20 or more servers becomes complex, increasing cost and reducing reliability due to human error. This is driving many companies to consider I/O consolidation/virtualization strategies that replace discrete Ethernet, FC (and sometimes InfiniBand [IB]) adapters with a pair of high-performance converged I/O adapters.

While typical server I/O bandwidth requirements are quite low, as more demanding workloads are virtualized I/O bandwidth requirements will increase. Virtual machine (VM) movement (VMotion) and backup traffic can drive requirements beyond 1 Gbps, and virtualization increases the required bandwidth significantly. Table 1 contains Gartner's projections for I/O bandwidth

requirements in early 2011. These projections are the result of Gartner research and they have been validated with a number of server, storage and networking vendors.

Table 1. The Impact of Virtualization on Server I/O Bandwidth Requirements*

Server Type	OS Instances	FC B/W per Instance	FC B/W per Server	TCP B/W per Instance	TCP B/W per Server	B/W per Server	Servers per Rack		Gbps per Rack
1 RU									
No Virtualization	1	200	200	100	100	300	36		**11**
With Virtualization	10	200	2,000	100	1,000	3,000	36		**108**
Blade Servers							**Servers per Chassis**	**Gbps per Chassis**	**Gbps per Rack**
16 Blades no VMs	1	200	200	100	100	300	16	**5**	**14**
16 blades with VMs	10	200	2,000	100	1,000	3,000	16	**48**	**144**

B/W = bandwidth, FC = Fibre Channel, I/O = input/output, OS = operating system, RU = rack unit, VM = virtual machine
*Unless otherwise noted, all bandwidth is in Mbps
Bandwidth estimates reflect anticipated average requirements in 2010 to 2011

Source: Gartner (August 2009)

With minimal consolidation of 10 VMs per physical server, virtualized 1RU or 2RU servers will likely require as much as 3 Gbps each, and blade systems will easily generate 48 Gbps of I/O per chassis. In order to minimize the impact of congestion-induced latency, the degree of acceptable overbooking in the aggregation and core network tiers will have to be reduced.

While FCoE is the most visible option, it is not the only approach and not necessarily the best (see "I/O Is the New Frontier of x86 Virtualization," "Virtual I/O Solutions for Virtualized Servers," "The Folly of Fibre Channel Over Ethernet," "Emerging Technology Analysis: Fibre Channel Over Ethernet, Networking and Communications" and "Case Study: Cloud Provider Achieves Ten Times Lower Total Cost of Ownership Over FCoE With Input/Output Virtualization").

Inhibitors to Backbone Convergence

If convergence within the server and rack is a good idea, why shouldn't you extend it across the backbone? Important factors include the following financial and technical risks:

- Gartner estimates that there is a $50 billion installed base of business-critical FC attached storage.
- In many configurations, a converged network requires more switches and ports, is more complex to manage, and consumes more power and cooling than two well designed separate networks (this will be covered in a forthcoming Gartner document titled "The Myth of Fewer Ports — FCOE and Data Center Network Convergence").
- CEE depends upon a handful of standards that are being developed by three different standards bodies, with completion of the suite not expected before late 2010.

A Three-Step Program to the Ideal Top-of-Rack Converged Network

Step 1: Consolidate and Virtualize the Ethernet I/O

As detailed above, virtualized servers require as many as six to eight 1 Gbps Ethernet network ports and two FC ports. In non-virtualized servers, three to four Ethernet ports are typical. The most conservative approach is to start by consolidating the Ethernet I/O and the associated switch ports. While Ethernet/FCoE may be the most visible approach, this can be accomplished via IB or native 10 Gigabit Ethernet (10 GigE), and in the near future via Peripheral Component Interconnect Express (PCIe)-based approaches.

Top-of-Rack Topology

Because of the distance limitations on copper cabling, and increasingly due to increased bandwidth density requirements, most deployments of Ethernet I/O virtualization use a top-of-rack switch/gateway, rather than an end-of-row switch to aggregate traffic (a follow-up report will address cabling issues). Many organizations will choose to maintain separate IB connections for compute clusters and FC connections for SANs (see Figure 1). This is the least disruptive method, and offers the lowest-risk approach to adopting I/O consolidation. Some organizations, primarily small to midsize enterprises, may take an additional step and consolidate storage traffic on Ethernet via Internet Small Computer System Interface (iSCSI).

Figure 1. Discrete SAN and LAN I/O

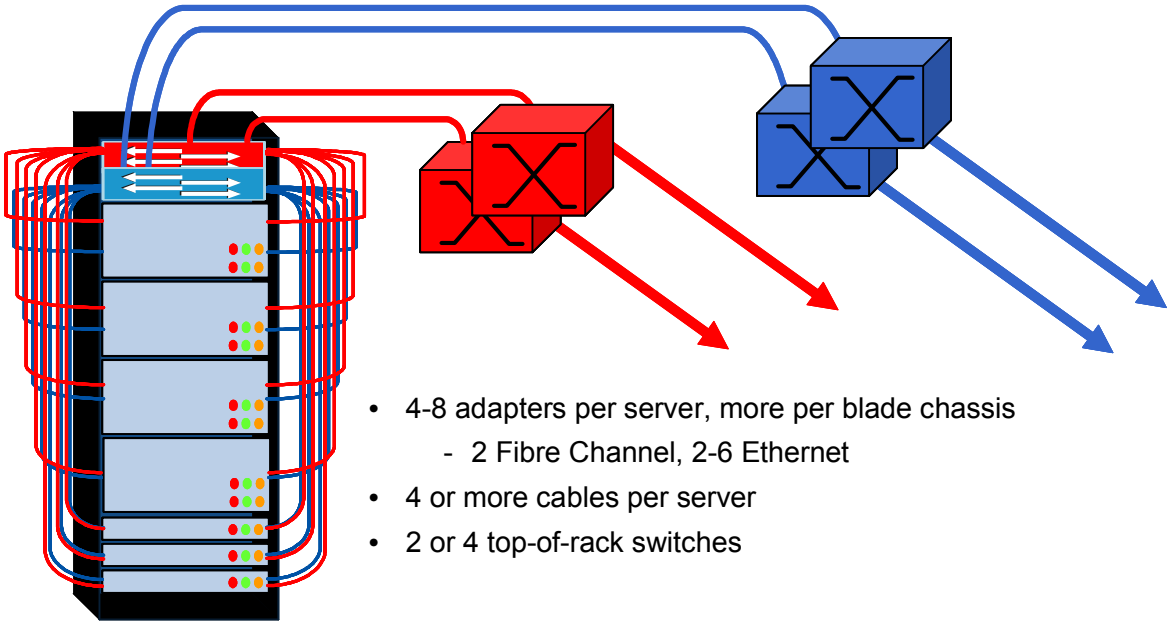

I/O = input/output, SAN = storage area network

Source: Gartner (August 2009)

Step 2: Consolidate and Virtualize all the I/O — The 80/20 Rule

For new build-outs of general-purpose servers, some organizations may choose to consolidate all server I/O. Using this approach further reduces cable, adapter and top-of-rack switch count by two-thirds or more. This can save easily 500 to 1,000 watts per rack. While FCoE is emerging as a preferred technology for this approach, a number of organizations have chosen to adopt IB-based approaches (see Figure 2).

Figure 2. FCoE in Rack Convergence With Top-of-Rack Switches

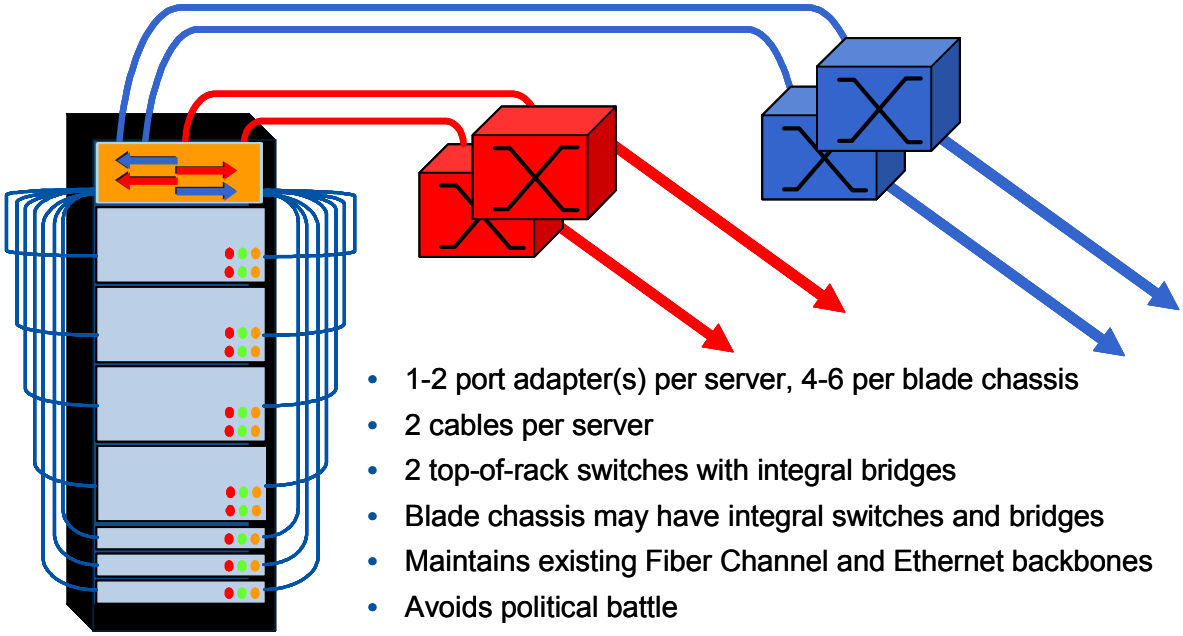

FCoE = Fibre Channel over Ethernet

Source: Gartner (August 2009)

For FCoE, this configuration of the top-of-rack switches includes an FCoE to FC bridge, allowing the FC and Ethernet traffic to be separated and sent to the respective existing backbone networks. For IB, an IB-FC and IB-Ethernet gateway is integrated with the switch. Connections to the existing SAN and Ethernet backbones preserve the investment in existing switches and management tools and eliminate the political battle over who controls the backbone switch. It also allows time for the remaining CEE standards to be finalized and for vendors to prove interoperability. When the standards are complete and interoperable products are available, the next phase can be considered.

Adopting this approach today, and during the next 12 months would be risky, because existing FCoE to FC bridges may lack adequate performance to forward server traffic to existing FC-connected storage arrays.

Some vendors suggest that storage be distributed across the data center into server pods and connect directly to FCoE top-of-rack switches. This approach limits storage aggregation and the dynamic placement of server workloads, negating the advantages of server and storage virtualization while driving up cost and complexity.

A better approach is to delay full convergence and to use FCoE converged network adapters (CNAs) for Ethernet I/O consolidation. Additionally, only acquire CNAs from vendors that guarantee that their existing CNAs fully comply with the recently ratified T11 standards, and that they interoperate with multiple FCoE switch vendors, most notably Blade Network Technologies, Brocade and Cisco.

Step 3: Plan to Move All Data Center Traffic to a Single Technology, Not Necessarily a Single Network

There is a benefit in standardizing on a single technology for all data center networking, if that technology adequately supports the needs of applications. Traditionally, the technology that adapts to assimilate others has been Ethernet. If CEE proves a suitable substitute for FC, it will make sense to eventually move all traffic to Ethernet technology. It will simplify acquisition, training and sparing. However, settling on a single technology does not require that the networks be combined. Design, operations and troubleshooting may be much easier with two separate networks. It may also cost less to build two separate networks.

The Payoff

In our model we assume that 10 Gbps CEE top-of-rack switch ports will cost approximately the same as normal 10 GigE ports. In fact, we expect CEE to become standard on most top-of-rack form-factor switches. Additionally, our backbone model shows that with the emerging generation of CEE backbone switches there is no saving in terms of total data center backbone switch ports. Consequently, the total savings occur within the servers and racks.

So, this approach delivers 80% (or more) of the benefits of a converged data center network, with only 20% of the work of a fully converged network. It reduces power and cooling problems in the rack. It preserves the investment in existing switches and storage arrays, and avoids the potential political fallout from trying to merge SAN, server and LAN teams.

It may also simplify the upgrade path to higher-speed uplinks (40 or 100 GigE) if the top-of-rack switches you choose have open expansion slots for 40 Gbps uplinks. However, if these slots have been used for FCoE to FC bridges and FC ports, this upgrade will be more disruptive because it will require a simultaneous switch to native FCoE storage connections.

The promise that unified fabric will require fewer switches and ports, resulting in a simpler network that consumes less power and cooling, may go unfulfilled. But that doesn't mean that enterprises should forgo the benefits of I/O consolidation and virtualization and a unified network technology. In fact, this 80% approach may prove that less is more (than the 100% solution).

RECOMMENDED READING

"I/O Is the New Frontier of x86 Virtualization"

"Virtual I/O Solutions for Virtualized Servers"

"The Folly of Fibre Channel Over Ethernet"

"Emerging Technology Analysis: Fibre Channel Over Ethernet, Networking and Communications"

"Case Study: Cloud Provider Achieves Ten Times Lower Total Cost of Ownership Over FCoE With Input/Output Virtualization"

REGIONAL HEADQUARTERS

Corporate Headquarters
56 Top Gallant Road
Stamford, CT 06902-7700
U.S.A.
+1 203 964 0096

European Headquarters
Tamesis
The Glanty
Egham
Surrey, TW20 9AW
UNITED KINGDOM
+44 1784 431611

Asia/Pacific Headquarters
Gartner Australasia Pty. Ltd.
Level 9, 141 Walker Street
North Sydney
New South Wales 2060
AUSTRALIA
+61 2 9459 4600

Japan Headquarters
Gartner Japan Ltd.
Aobadai Hills, 6F
7-7, Aobadai, 4-chome
Meguro-ku, Tokyo 153-0042
JAPAN
+81 3 3481 3670

Latin America Headquarters
Gartner do Brazil
Av. das Nações Unidas, 12551
9° andar—World Trade Center
04578-903—São Paulo SP
BRAZIL
+55 11 3443 1509

Gartner