

The corporate network of the future: convergence 2.0

Date: *Fri, 10/29/2010 - 13:45* Source: By Alberto Soto, vice president, EMEA at Brocade

How convergence and the cloud will simplify network operations and drive business performance, but it won't happen overnight



Alberto Soto, vice president, EMEA at Brocade, at NetEvents EMEA Press Summit Istanbul

What do CIOs and their teams want from the corporate networking technologies of the future? Lower latency, enhanced reliability and higher throughput are all high on the IT department's wish-list – but, above all, they say, what they'd really like to see is a single network infrastructure capable of handling all data, clustering and storage traffic.

It sounds like a tall order, but it's hardly surprising that many CIOs are taking a long, hard look at their current network infrastructures and finding them wanting. Today, their teams are typically managing two or three parallel networks: they have a storage network built for reliability, guaranteed data integrity and non-blocking performance. And they have a data network whose characteristics include best-effort performance and unpredictable bandwidth, along with often-frustrating levels of complexity. Separate switches, host bus adapters (HBAs), network interface cards (NICs) and cables are required for each network.

This clutter imposes a significant burden on IT department time and budgets and has environmental consequences too, as network administrators struggle to provide each network component while managing the power and cooling each requires.

The race to consolidate servers and storage systems through virtualisation, while delivering considerable efficiencies, has only added to the demands imposed on the corporate network infrastructure. By allowing multiple applications to run on a single server, IT teams have been able to cut the number of physical servers required, slash the energy consumed and vastly increase operational efficiency.

However, a single server hosting perhaps 20 virtual machines (VMs) requires significant bandwidth to keep these services up and running. The dynamic movement of VMs between different physical machines, according to their processing requirements and users' needs, poses further network-management challenges. If network connections don't travel smoothly with their associated VMs, or create unpredicted demands on physical servers during times

of peak processing, network performance - and the end-user experience - quickly goes downhill.

The next steps along this journey will present further challenges for network administrators. Increasingly, companies are capitalising on their early wins with virtualisation to create internal (or private) clouds, where all IT assets interact seamlessly to provide a single 'pool' of computing resource. Beyond that, they plan to work towards a 'hybrid' cloud model, where applications and services are provided by a mix of internal and external, third-party resources. Either way, the cloud – where distributed applications, residing on different machines, all perform different parts of a vital business process – clearly means more traffic traveling across corporate networks and beyond the firewall.

As a result, CIOs expect tomorrow's corporate networks to fulfill a wide range of sometimes-conflicting demands. They want unprecedented scalability, but reduced management complexity. They want seamless mobility, but tight orchestration and they want emerging networking technologies to complement the investments they are making today, instead of forcing them to refresh the entire environment in a wholesale 'rip-and-replace' exercise. For these reasons, unified fabrics are creeping up the CIO's agenda.

The promise of convergence

Tomorrow's networking environment will consolidate user-application traffic and storage-data traffic onto a single, high-performance, highly available network that has the built-in intelligence to identify different traffic types and handle them appropriately, according to pre-defined rules. The benefits of the concept are clear in terms of time, cost, skills and procurement, not to mention reduced cabling complexity. And although the path to getting there may at first seem steep, an incremental approach can ensure that the journey needn't cause unnecessary upheaval for the IT team and the end-users that they serve.

We are in the midst of a major innovation cycle in IT – moving from infrastructure designed to automate existing processes to services-oriented architectures that speeds the delivery of innovative new services. Adopting new technologies means enabling new business models and transforming relationships with customers, partners and employees.

The days when the majority of computing power was in the data centre are behind us. Today, we have incredibly smart end points with lots of computing power that are remote, distributed and mobile. Information and applications are virtualized and can reside anywhere within the infrastructure we refer to as the cloud.

The compute model has been reversed – both information and the consumption of that information are distributed. This shift dramatically drives the importance of the communications throughout the network. And because you have distributed the compute power and your information storage, you have in essence distributed the data centre.

As data centres become distributed, the network infrastructure must take on the characteristics of a data centre. And if the network becomes our data centre then the network is our business. In this highly distributed and dynamic environment, your commercial model, your customer experience and your employee productivity are all reliant on just how robust your network is.

A very established and well-understood technology lies at the heart of the network - Ethernet. It is already the predominant network choice for connecting servers to each other

on the corporate local area network (LAN) for the purpose of transporting user-application traffic; the convergence 2.0 concept proposes that the separate Fibre Channel network, that typically transports storage-data traffic around a storage area network (SAN), be shifted to Ethernet, too – or Data Centre Bridging (DCB).

But in order for this kind of SAN/LAN convergence to work, DCB must offer the reliability and latency characteristics associated with Fibre Channel technology. That's because storage networks require data to be delivered in sequence and intact - in industry parlance, a "lossless" infrastructure is required. Ethernet, by contrast, falls short in this respect, taking instead a 'best effort' approach, where data is not necessarily delivered in the right order and where some packets may be dropped altogether due to network congestion.

In response to these challenges, the essence of DCB lies in higher transmission speeds, based on 10-Gigabit Ethernet (GbE) technologies, and enhancements to the underlying Ethernet specifications in order to replicate the reliability and class-of-service features seen in today's SAN environments.

Today, standards bodies representing some of the world's foremost suppliers of storage and networking technologies are working on specifications to increase the performance of existing Ethernet networks and to tunnel Fibre Channel traffic securely and efficiently through Ethernet infrastructures.

One of the most important of these is the Fibre Channel over Ethernet (or FCoE) standard, an encapsulation protocol that wraps Fibre Channel storage data into Ethernet frames, enabling it to be transported over a new lossless Ethernet medium. Developed by the T11 technical committee of the International Committee for Information Technology Standards (INCITS), it relies on flow control to recognise when a buffer is almost full and to request that the sender stop transmission until the buffer has emptied and the transmission can start again.

One of its major advantages is that FCoE will use the same FC drivers, switches, cabling and management applications already in use today. That is allowing companies to start introducing concepts such as unified fabric within the first few metres of their data centre networks today, using FCoE with DCB at 10Gbps speeds to send both data and storage traffic to the first access switch it encounters. That traffic can then diverge to the corporate LAN or to the existing SAN using Fibre Channel, accordingly, and the number of cables that need to be plugged into a server rack are immediately reduced.

The road ahead

In this way, a technology still in its infancy is already starting to take shape at the network edge in the data centres of early adopters. Others will follow their example. Richard Villars, vice president of storage systems at market analyst firm IDC has predicted that 2010 will see an increase in converged networking pilot projects, with significant technology deployments expected in 2011. According to figures from analyst firm Dell'Oro Group, approximately 10,000 FCoE ports shipped in 2008, but the company's analysts anticipate that this figure will rise to about 1 million in 2011.

Of course, mainstream convergence is dependent on other factors, too. Storage device vendors must allow for FCoE adapters on their products. Network interface cards (NICs) for 10GbE must continue to drop in price and be incorporated into server motherboards. And

within end-user organisations themselves, CIOs will need to reorganise. They may wish to ensure that storage, server and network operations can be monitored by a single service desk, for example. Network and storage teams will need to work together far more closely than they do in today's silos. And when data centre switches are ready for replacement, they will need to ensure that whatever new solutions are bought can support the move to unified fabric.

But reporting on their findings, Forrester analysts described the use of Ethernet for both LAN and SAN as a concept with "significant momentum", and one that is likely to provide considerable benefits for adopters. "While there remain some questions about what it will look like and when most firms will move towards adoption," they added, "it makes good sense to gain further understanding of the concepts and supporting technologies, and to begin evaluating Ethernet storage offerings now as a precursor to eventual fabric unification."

Some in the industry would have end users believe that the converging of storage and data networks is a relatively easy process that should be done as soon as possible. But this assertion must be challenged as businesses will migrate in their own timeframe. The complexity of managing and deploying a network and data centre is reaching unacceptable levels.

With the advent of the Cloud and the drive to convergence underway, this paradigm shift is enabling the ability to selectively outsource or extend a significant portion of the IT infrastructure. Now businesses can adopt a wide range of models depending on the needs of the business. Some may want to move rapidly to a "pay-as-you-go" model, while others may migrate slowly to a private model built internally.

Either way businesses are taking advantage of the convergence revolution by adopting new operating models and creating virtual enterprises. In doing so, they become much more agile and flexible, more efficient and less capital intensive and becoming more competitive but only if the architecture is affordable and reliable.

1 Both reports indicate strong interest among respondents in LAN/SAN convergence: 'Benefits Of SAN/LAN Convergence', published by Forrester Consulting, December 2009; 'Future Datacenter Management Strength – Network Convergence', by Reco Li & Jason Chen, published by IDC, October 2009

2 'Future Datacenter Management Strength – Network Convergence', by Reco Li & Jason Chen, published by IDC, October 2009

3 'SAN 5-Year Forecast Report', published by the Dell'Oro Group, August 2009

4 'Benefits Of SAN/LAN Convergence', published by Forrester Consulting, December 2009

<http://www.telecomkh.com/en/business-communications/brocade/cloud-computing/convergence-20/2441>