CLABBY ANALYTICS

Case Study

Network Strategic Planning for Next Generation Data Centers

Introduction

Major changes are underway in the networking industry that will have a major influence on the way that data centers are designed and managed. Two of the most prominent "macro trends" that we see will be in the areas of:

1. Converging network fabrics over Ethernet
2. Network virtualization and management of Virtual Machines (VMs)

In the areas of converging network fabrics over Ethernet, standards committees are close to finalizing both the Fibre Channel over Ethernet (FCoE) and Data Center Bridging (DCB) standards as a means to converge Fibre Channel Storage and LAN traffic over the same Ethernet network. This convergence will simplify network design (reducing the number of network components needed to connect different network topologies with each other — as well as simplifying management by using a common approach to networking across entire systems and storage tiers).

In the area of managing Virtual Machines, there is an argument that suggests that *network devices may be better suited to host that logic*. In order for network switches to fully support virtualization, network switches need to be aware of Virtual Machine activities, locations, image requirements and rights/privileges — rather than just physical server connections. Hosting this logic and intelligence with a network switch or across a network fabric could serve to reduce virtualization handling on CPUs — thus freeing up CPUs to do what they do best: application and database processing.

Enterprises are expected to handle increasing demands, yet with reduced budgets and resources. Large enterprises are looking for ways to save on operating costs and avoid making new capital expenditures, and conversely, want a cost-effective way to scale-out or scale-up. The data center itself is also the focus of company cost-cutting efforts — companies ask what can be done to make better use of the investments they have already made? What new virtualization technologies should they be evaluating? How can they reduce the complexity of the data center by consolidating networks, servers and storage – thereby making management simpler and driving down power and cooling costs? And, large enterprises need to accomplish this without sacrificing performance. Further, enterprise customers want to be sure their data center will support future growth and related bandwidth requirements.

In this *Research Report*, *Clabby Analytics* examines these trends and their effect on helping data centers achieve their reduce-cost-and-complexity needs, while handling increasing demand within limited budgets and resources. We use BLADE Network Technologies, a rapidly growing, industry-leading data center networking solutions provider, to illustrate how one vendor is strategically addressing the converged Ethernet fabric and network-based

virtualization trends — and to highlight how the competitive landscape will change vis-a-vis product designs, strategic partnerships, and in specialized, industry-focused solutions.

A Closer Look at How These Trends Will Affect Future Data Center Designs

We've all heard about cloud computing — public clouds, private clouds, hybrid clouds — but a private cloud is really just the next iteration of the virtualized data center. The next generation data center will consolidate and virtualize compute, network and storage as a means to simplify resource management while increasing resource utilization. The trend toward converged networking aims to address the simplification of resource management; the trend toward virtualization at the network level aims to improve resource utilization. This section takes a closer look at each trends — and the benefits that can be derived by converging network fabrics and moving some of the control of virtualized resources to the network hub/router level through switch-resident software.

A Closer Look at the Network Convergence Trend

One of the biggest problems in networking to date has been the use of differing wiring schemes and protocols to allow various systems and storage devices to connect and communicate with one another. This had led to increasingly complex data centers with a mix of adapters, cables and switches that are costly to maintain and manage. Different "pipes" for connecting devices now include InfiniBand, Fibre Channel, variable speed Ethernet and more.

Ethernet dominates — but other approaches are still being used, and these approaches need to be accommodated (often, by using gateways). The big driver behind the use of these differing wiring schemes and associated protocols has been the need for optimal bandwidth at various tiers within an organization. Data Center Bridging (DCB), however, remedies the need for differing wiring schemes by providing a single standard for connectivity and communications.

Data Center Bridging, also known as Converged Enhanced Ethernet (CEE), is an enhancement to the Ethernet protocol to prevent data loss without affecting performance. DCB's enhancements to standard Ethernet include:

- No packet loss
- Priority-based flow control
- Support for low-latency connectivity requirements and,
- Congestion notification and management.

Thus, while network convergence means fewer cables, adapters and switches - reducing cost and complexity and easing management, it requires two things: 10 Gigabit Ethernet (10GbE) transmission speed and DCB.

Standards committees are close to finalizing both the Fibre Channel over Ethernet and Data Center Bridging standards as a means to converge Fibre Channel Storage and LAN traffic over the same Ethernet network. This is beneficial since each of these networking technologies has its advantages. Ethernet has cost advantages over Fibre Channel and is widely deployed, and, because of Ethernet's widespread market acceptance, plans already exist for higher performance (40 Gb/s and 100Gb/s) networks.
But, Ethernet also had a downside: it is not typically a "lossless" network like Fibre Channel, since the Ethernet protocol may lose data packets if the network is congested, making it less reliable than Fibre Channel. However, DCB makes Ethernet lossless and as reliable as Fibre Channel.

By leveraging DCB's Ethernet enhancements and network convergence, FCoE can use Fibre Channel's network, service and protocol layers to transport data over Ethernet's physical and data link layers. And, with the convergence of storage, networking and clustering data into a single fabric, application performance is increased while cutting energy and cooling costs, reducing the required number of cables and network adapter cards and simplifying management. Applications do not have to be modified to run over FCoE. With the FCoE and DCB standards predicted to be finalized in the coming year, IDC predicts 2010 will see an increase in converged networking pilot projects with significant technology deployments expected in 2011 (source: "Converged Network Fabrics Reduce Complexity, Cabling Installation and Maintenance," November 1, 2009). DCB (CEE) also provides benefits for iSCSI and NAS, as much as it helps FcoE. Thus, DCB brings many benefits to data centers that converge their fabrics on Ethernet.

Clabby Analytics believes that as these standards become finalized, Ethernet will emerge as the dominant data center networking protocol and continue to be in the future as Ethernet moves to 40G and 100G speeds.

Network Virtualization is Necessary to Realize the Benefits of Server Virtualization

To take advantage of server virtualization, data centers need to virtualize their networks. In order for network switches to perform virtualization, network switches need to be aware of Virtual Machines' activities, locations, image requirements, and rights/privileges — rather than just physical server connections.

Only BLADE's Virtual Machine-aware switch-resident software, VMready, is capable of addressing this need for easy identification and automatic management of Virtual Machines' network policies as they move. VMready is a key element of BLADE's Unified FabricArchitecture, which we describe in more detail below.

One topic of discussion that has long been discussed by strategic planners is "where should the logic reside for handling virtualization tasks". Indeed, systems can run virtualization logic — but storage and networks devices are also capable of running virtualization logic. And in many cases - network, storage and server virtualization can be complementary, offering combined benefits when implemented together.

The Key to Exploiting These Trends: Look Beyond Networking Features/Functions to Cloud Service/Support; Strategic Partnerships; and Industry-specific Solutions

All communications/networking vendors look alike — or do they? A closer look at this industry shows very significant differences in communications/networking hardware and software features, function, performance and reliability. In addition, different service and support strategies are emerging, new strategic partnerships are developing, and data center networking solutions are evolving to meet, for example, the low latency demands of Wall Street, or the fast performance for HPC applications such as Web 2.0, IPTV, video-on-demand, oil and gas exploration and cloud computing.

The vendor that we have chosen to illustrate the kinds of leading-edge networking solutions needed to address next-generation data centers is BLADE Network Technologies (BLADE for short). The following section provides an overview of BLADE, followed by a discussion of how converged networking over Ethernet fabrics and virtualization-at-the-network-layer have been blended into BLADE's overall strategy.

BLADE Background

BLADE Network Technologies, founded in 2006, is a privately-held company headquartered in Santa Clara, California, with global presence. It is:

- A leading provider of data center Gigabit and 10G Ethernet networking solutions
- A market leader in converged Ethernet fabrics
- An innovator in network virtualization and management software

The company's products include: 1 and 10GbE switches for blade servers, top-of-rack switches (BLADE just added a DC-powered RackSwitch for scale-out containerized computing to its line-up), as well as virtualization and rack management infrastructure software. The BLADE switch family is:

- Lossless
- Low latency
- Low power
- Low cost
- Open, standards-based and interoperable with other vendors, such as Cisco

BLADE's Ethernet switch family boasts impressive growth statistics including 54% unit growth in 2008, 25% growth from 2008 to 2009 (in a down economy), over 50% market share in blade switches, and well over 8 million ports deployed. Dell'Oro Group (source: "Ethernet Switch Report", Dell'Oro Group, November 2009) has named BLADE #3 in Gigabit Ethernet. BLADE's impressive growth is partially due to its strong product offerings, but it can also be attributed to the assistance of strong partners including IBM, HP, NEC, Netezza and SGI, and an ecosystem that includes technology partners such as Broadcom, Emulex, Intel, Microsoft, VMware, Citrix, Chelsio and NetApp,

> ***Note: BLADE's rapid market growth has come at a time when many technology companies are struggling to remain profitable — a testimony to BLADE's market leadership in Ethernet-based data center networking, and its strong track record in helping customers achieve optimal TCO while improving performance.***

What Differentiates BLADE Network Technologies' Products

BLADE competes primarily against Cisco Systems, Brocade and Force10 Networks.

Performance

From a performance/power consumption perspective, an independent study conducted by the Tolly Group, (source: "BLADE Network Technologies RackSwitch G8100 Series: Competitive Performance Evaluation versus Cisco Catalyst 4900M Switch", Tolly Group May 2009) provides an example of how BLADE has created a performance edge in the switching marketplace. The findings showed that BLADE's switch offered:

- *8.2x less latency* for various frame sizes in a port-to-port configuration;
- *2x higher throughput* for all frame sizes tested up to 9,216 bytes in a 24-port full mesh test;
- *Consumed 65% less energy*; which translates to a savings of $198.37 per year, per switch deployed
- *5.35x better price/performance* on average.

To create performance and power consumption advantages, BLADE uses a unique approach to scaling network bandwidth for consolidated rack/blade server/storage environments. The company calls this approach "Rackonomics." Rather than tying racks and blades into expensive, external core switches, BLADE has architected switch modules that can easily fit into existing blade server/rack environments — enabling networking bandwidth to be scaled cost effectively within a rack, row or container. BLADE's "faster, virtual, proven" Unified FabricArchitecture provides significant reduction in data infrastructure total-cost-of-ownership (TCO). BLADE's 3 million hour MTBF means BLADE offers high reliability as well price/performance and lower power consumption advantages with its lossless, low latency, low power and low cost Ethernet switches.

Rackonomics

Rackonomics is an approach that reduces cost and complexity by duplicating pre-provisioned racks, rows of racks or containers. Rackonomics enables IT managers to save tens of thousands of dollars in switch hardware, deployment, and energy costs for every external core switch they are able to avoid deploying.

Using the BLADE Rackonomics approach, data center architects can standardize on a unified and affordable rack-level network infrastructure to cost-effectively provision and scale-out Web 2.0 environments, high-performance computing (HPC) clusters and virtualized data centers.

Innovation Air Flow/Heat Dissipation

To further differentiate its hardware offerings, BLADE's switches have two distinct advantages in airflow/heat dissipation as compared with both traditional rack switch designs as well as to core switches. These advantages are manifest in a unique, patent-pending airflow design; and in advantages that can be gained by exploiting cooling at the rack or blade level. With respect to airflow/heat dissipation, BLADE has designed an innovative new approach to circulating air over its rack switches. Most rack switches have been designed with airflow that flows in the opposite direction of server airflow— generating a condition known as "hot loops." These hot loops actually increase cooling requirements by drawing warm air over already hot blade or rack servers. BLADE's Ethernet switches are designed to draw cool air from the front of the rack or blade cabinet— enabling front-to-back airflow that allows for cool aisles and hot aisles (a more energy efficient approach to heat dissipation).

BLADE's Unifying Data Center Framework: "Unified FabricArchitecture (UFA)"

BLADE' describes UFA as a "faster, virtual and proven" data center networking platform that allows customers to tie together best-of-breed servers, storage, hypervisors and core networking systems. Simply put - UFA is a converged network solution that allows SAN, NAS and HPC clusters to run over a single 10G Ethernet fabric.

The following subsections look more closely at the key elements of BLADE's Unified FabricArchitecture.

BLADEOS

BLADEOS is the world's most widely deployed switch operating system with well over 8 million ports in production deployment. BLADEOS is the foundation for a number of unique extensions that differentiate BLADE from other networking vendors. One is a converged fabric extension to BLADEOS that allows customers to converge their LAN, SAN and management networks onto a single wire. Another is vNIC which provides the ability to carve up a physical NIC into multiple virtual NICs and create a virtual pipe and allocate bandwidth

between the adapter and the switch port. Both of these capabilities will help shape the future of datacenter networking, being open and designed to support everything from mainframes to rack optimized servers to tower systems to purpose built application appliances (e.g. data warehousing, security, etc.) from multiple vendors.

BLADE's Network Virtualization Approach

One of the cornerstones of BLADE UFA is VMready, BLADE's switch resident network virtualization software that enables networks to be Virtual Machine (VM) aware. Many network switches do not have the intelligence to be VM aware. When IT moves VMs manually, the network policies associated with those VMs may be incorrectly configured with the potential for outages and security breaches. It is only when the network is VM aware that the benefits of server, network and storage virtualization can be fully realized. VMready assigns network policies on a VM level to guarantee the appropriate allocation of networking resources - and monitors and detects events in the virtual environment so that network policies remain associated with their virtual machines no matter where they reside.

An open architecture, VMready works with all virtualization products, including VMware, Microsoft's Hyper-V, Xen, Oracle VM and KVM, without any modifications.

> ***BLADE's open architecture approach prevents vendor lock-in, and enables companies to protect their existing investments. And, since VMready is switch-resident, it is not using CPU cycles which are required for application processing.***

VMready automatically synchronizes with VMware vCenter to create port groups. These port groups all have the same network configuration on all the required ESX vSwitches. This automatic configuration simplifies administration and reduces configuration errors.

VMready also tracks the mobility of virtual machines across the data center and automatically reconfigures the network in real-time as the virtual machines move. Consistent network policies are enforced regardless of a virtual machine's physical location, simplifying management and reducing operational expenses. By balancing resources to the application workload, BLADE's NMotion provides the highest application performance and availability.

The FabricHarmony Suite includes BLADEHarmony Manager, Smart Server Control (configuration management) and Open Fabric Manager (data center fabric management). BLADEHarmony Manager provides remote monitoring and management of BLADE switch modules, offering centralized administration for large groups of switches including auto-discovery, automated software downloads, backup and restore, cloning of configurations, event notification, health status reporting, performance monitoring. It extends VMready by tracking and reporting virtual machine network location/movement and provides NMotion support between independent (unstacked) VMready switches across the data center. Large scale deployment/management of VMready technology is provided by centralized, synch-ronized configuration of VM network policies in switches and hypervisors and port group configurations in the ESX vSwitches. BLADEHarmony Manager can centrally define policies — then automatically distribute configuration information to the switches and hypervisors — simplifying management. BLADEHarmony Manager integrates with existing management systems including HP Systems Insight Manager, IBM Systems Director and VMware's vCenter.

Smart Server Control is the industry's most sophisticated server control software used by over one million blade servers deployed in data centers around the world. Smart Server Control is a configuration management tool providing environmental, configuration, I/O and provisioning control.

Open Fabric Manager is designed to help manage growth and complexity by making it easy to manage storage I/O and network interconnects for up to 100 IBM BladeCenter chassis and up to 1400 servers. Open Fabric Manager virtualizes network parameters such as the World Wide Name (WWN) and Media Access Control (MAC) addresses. When you replace a blade or failover from one blade in a chassis to another blade in a different chassis, the LAN and SAN configurations are not affected. Once installed, the utility is resident in the built-in Advanced Management Module (AMM) so you can pre-configure LAN and SAN connections — I/O connections are made automatically when you plug in a blade. This is easily managed through a Web-based user interface based on IBM's Systems Director 6.1. An advanced version adds the capability for automated I/O failover to standby blades.

BLADE in Converged Networking: Strong Products Complemented With Strategic Partnerships
In short, BLADE builds excellent 1/10Gb Ethernet switches — but to broaden its product reach into converged networking; it has proven to be more expeditious for BLADE to partner with industry leaders rather than trying to build entirely converged networks on its own. To complete its UFA framework and broaden its range of solutions, BLADE has structured strategic relationships with companies such as Netezza, SGI, NetApp, Chelsio and Emulex.

- BLADE's strategy is to remain focused on data center networking, rack infrastructure management and Virtual Machine-aware network virtualization solutions. BLADE maintains strong relationships with server vendors including IBM, HP and NEC — who control over 80% of the blade server market.
- In a recent partnership announcement, BLADE and Emulex announced the delivery of networking components for IBM's BladeCenter Virtual Fabric. Using BLADE's BNT Virtual Fabric 10G Switch modules with vNIC and the Emulex Virtual Fabric Adapter CFFh for IBM BladeCenter provides an end-to-end virtual network fabric for the IBM BladeCenter.

This combined solution provides a virtual fabric that:

  - Carves up virtual pipes between the adapter and the switch, which provides maximum performance per virtual port. In addition by using virtual pipes you actually improve security and availability by isolation vNIC communications and reducing the potential of a failure in one vNIC, virtual group or uplink impacting the others.
  - Allows a dual port 10Gb adapter to be carved up into as many as eight virtual ports (4 per port). Clients than have the ability to assign bandwidth in increments of 100Mb up to a total of 10Gb per port.
  - Use 50-75% less hardware (adapters/switches);
  - Reduces cable clutter by up to 58%;
  - Manages 4 virtual groups instead of 112 virtual NICs
  - 50% more energy efficient

> ***This joint solution from IBM, BLADE and Emulex today allows for carving up Ethernet pipes using industry standard technologies and will also provide DCB/FCoE support delivering BLADE customers with: (1) hgh performance lossless storage over 10GE; (2) a single adapter for both FC and Ethernet traffic; (3) lower server costs because of fewer i/o slots; and, (4) lower management cost because of unified fabric. BLADE is also working closely for a solution leveraging iSCSI with HW offload which will also be extremely attractive to those looking at iSCSI deployments.***

At the end of 2009, BLADE delivered new network convergence capabilities for the BNT Virtual Fabric 10G Switch Module with support for FCoE using DCB. This specific solution delivered the industry's first and only fully integrated Fibre Channel over Ethernet (FCoE) solution inside a blade chassis. This integration significantly reduces power, cost, space and complexity over external FCoE implementations. The new solution converts CEE (Converged Enhanced Ethernet) packets to native Fibre Channel, so IBM BladeCenter customers can directly attach to their existing storage targets without any impact to their SAN infrastructure. By comparison, this integrated FCoE solution is 3x less expensive, 2x more energy efficient and requires no cabling between the switch and the gateway compared with external top-of-rack gateway designs.

> ***The DCB/FCoE support did not stop there - with the new BLADEOS releases BLADE's BNT Virtual Fabric 10G Switch Module and the BLADE RackSwitch G8124 also support interoperability with other industry vendors gateway devices like the Cisco Nexus and Brocade offerings. BLADEOS is the most widely deployed network operating system with well over 8 million ports in production deployment.***

Netezza has chosen BLADE as their partner for their Twinfin data warehousing platform. By combining top-of-rack and embedded blade switching with the Netezza Twinfin appliance, customers get the fastest end-to-end data center fabric infrastructure for data warehousing and analytics, delivering orders of magnitude faster performance, fast load speeds of up to 2 terabtyes/hour and fast backup and restore rates of as high as 4 terabtyes/hour.

Another Differentiator: Industry Specific Solutions

As network vendors compete in the cloud their customers will be looking to their network product suppliers for any "special" design/service/support offerings that their vendor can provide that can service industry-specific needs. BLADE has been active in helping to architect industry-specific solutions — and will continue to invest in network differentiation by industry in the clouds of the future.

An Example: Financial Services

In financial services, BLADE has found that applications are extremely demanding, requiring ever-increasing performance, high availability and reliability and minimal latency. For a financial trading firm, losing a few seconds can have a huge financial impact and outages can be catastrophic. BLADE has demonstrated competitive advantage in latency (8x lower than comparable Cisco products), which is particularly important in financial applications that have strict latency requirements.

Further, with 10GE and converged network fabrics from the BLADE/IBM/Emulex partnership, financial institutions (and other industries as well) can take advantage of

existing infrastructure, save money on cabling and network adapters, and with high performance 10GE and with DCB, take advantage of a lossless converged LAN/SAN fabric that can be managed as a single entity. BLADE is the network switch provider to a premier financial institution with over $170 billion in assets under management with over 1500 retail branches and approximately 3000 ATM's managing critical data center assets.

Summary Observations:

In this report, *Clabby Analytics* used BLADE Network Technologies to illustrate how one networking vendor is planning to compete in data center environments of the future. Other networking vendors are also concentrating their efforts in these areas (for instance, Cisco has partnered with EMC to fill in its network virtualization gaps; and HP has acquired 3COM to round out its converged fabric product offerings). But BLADE is also using innovations in power use/heat dissipation, virtualization, and in-rack infrastructure management industry-focused designs to differentiate itself from its competitors.

> ***BLADE's United FabricArchitecture is the foundation upon which these advanced technologies (converged fabric, network virtualization logic) are built.***

From a converged fabric perspective, what is most exciting about BLADE is that, through the partnership with IBM, Emulex and others, the company now has products that provide the foundation for converged DCB networks in the future. For large enterprises in a broad set of industries including financial services, healthcare and retail, the benefits of converged Ethernet network fabrics include cost savings in power, cables, network adapters and management. And these "cloud-ready' solutions are entering the market at a time when large enterprises are giving serious consideration to implementing private and public clouds. BLADE's converged Ethernet switches with network virtualization will be a compelling cloud foundation technology with the promise of immediate benefits. BLADE's partnership with Netezza gives them a high-performance solution for data warehousing and analytics, a rapid growing segment of the market and one where large enterprises are increasingly focused.

From a virtualization-logic-in-the-network perspective, it is important to note that BLADE has been quite active in this space. The company now offers VMready, Smart Server Control and the FabricHarmony Suite — and we expect more products in this space over time.

The strategic direction of the network market as it moves to next generation data centers is well known. All of the major networking vendors are driving to build converged networks and to enable networks to host virtualization logic. BLADE is well prepared for this forthcoming competition as networking companies compete for market share.

Clabby Analytics
http://www.clabbyanalytics.com
Telephone: 001 (207) 846-6662