THE CLIPPER GROUP
Navigator™

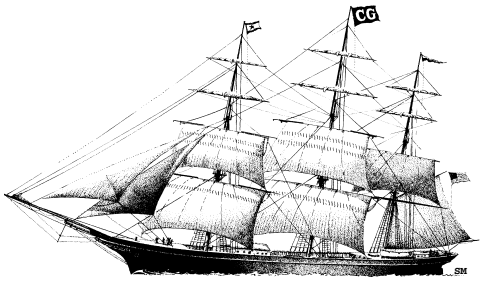

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2007097 November 13, 2007

IBM's Answer for Improved Data Center I/O — BladeCenter Open Fabric Manager

Analyst: David Reine

Management Summary

Down through the centuries, governments, and enterprises have been dealing with problems arising from the use of different rules and regulations from country to country. The formation of the European Economic Community is the latest attempt to set standards for use in commerce among European nations. However, the establishment of standards, in fact the need for standards in commerce, goes back several centuries. The effective implementation of a railway system in Europe owes its very existence to the establishment of a standard gauge for railroad tracks. The Gauge Act of 1846 in England established 4' 8 1/2 ", the existing gauge of horse drawn wagons, as the standard distance between the inner sides of the rails, and was later used to establish **a network** of railways throughout Europe, becoming known as "Standard Gauge". When steam-powered railroads appeared for the first time in North America, the early trains came from Britain, built to "standard gauge", and enabled commerce to flow smoothly between Canada and the U.S. The standard, however, predates nineteenth century England by many centuries. It also happens to be the distance between the wheels of the chariots of ancient Rome, which left many permanent ruts in the muddy roads and forced subsequent modes of transportation to conform. Two thousand years later, you can say that *we are still in the same rut!*

The requirement of standards does not stop at the railroad depot, however. For modern commerce to be successful, the enterprise data center must be able to communicate beyond the walls of the "glass house". The enterprise must be able to communicate with branch offices, distribution centers, and e-commerce partners around the globe. The servers that provide these communication lines must be able to talk in multiple protocols, across many different networks, and they have to be flexible enough to adapt to the constantly changing needs of the enterprise. The data center staff must be able to deploy a complete set of network adapters throughout the enterprise, supporting all of the features and functions of the switched network. Further, the data center must accomplish this deployment at low cost, with easy-to-use tools, and it must be accessible to all of the enterprise's ecosystem partners. Most importantly, it has to be easy to manage, as it may be installed in locations without an I.T. staff.

Today, the architecture of choice for both data center and remote office use is the blade, with IBM's *BladeCenter* leading the way. One reason may be the availability of IBM's *Open Fabric Manager* to manage a continually evolving communications network. To learn more about how IBM's Open Fabric Manager can connect your ecosystem, please read on.

IN THIS ISSUE

Today's Enterprise Data Center

In data centers around the world, Information Technology is rapidly approaching a crossroads. The uncontrolled proliferation of energy-dependent servers and storage has turned the acquisition, deployment, set-up, and management of the I.T. infrastructure into an expensive nightmare. In addition, power consumption is out-pacing our ability to generate affordable energy. The CIO can no longer support a scale-out environment where the data center deploys servers in a traditional application-specific fashion, with departments installing a new rack-mount server every time a new application needs to come on-line. There is enough evidence in front of us to prove that the average x86 server is only consuming 15-20% of the available compute cycles, but still consumes 100% of the power cycles. Every enterprise is doomed to run out of floor space, unless they run out of power first.

Total cost of ownership (TCO) has become the watchword in every boardroom, as management attempts to solve a never-ending stream of I.T. problems with a very limited I.T. budget. The enterprise needs to change its I.T. paradigm if it expects to achieve fiscal control. Consolidation of mission-critical enterprise applications on new low-cost, open systems servers with multi-core processors are proving to be an effective first step to gain control of energy consumption and floor space in the data center. By virtualizing the operating environment of the server through new hypervisor software such as *VMware*, *Xen*, and others, the data center can deploy multiple x86 applications on the same physical platform, increasing the utilization of server processing power to upwards of 75-80%, and reducing the server footprint by up to 80%.

Unfortunately, the TCO consists of more than just server acquisition costs, energy, and floor space. The data center also needs to support a communications network; consolidation of multiple applications onto a single platform does nothing to simplify that task. In fact, it may complicate network management further by forcing the I.T. staff to support multiple I/O adapters and multiple switches on the same platform, with the added burden of communicating from a single server to a heterogeneous communications network.

Consolidating with rack-mounted servers does little, however, to solve many of the issues of server proliferation, power consumption, and server footprint. Rack-mounted servers do even less in simplifying the administration and management of implementing and maintaining a complex communications network. Blade servers, however, do. They can reduce data center complexity and, therefore, reduce TCO for the data center infrastructure.

Managing a Consolidated Data Center

Bladed servers are not simply a new form factor for the same old server deployment. A blade chassis represents a significant change in the data center computing paradigm, a change for the better. They support more than just blade servers; they support storage and networking options, as well as a variety of commodity servers based on Intel's *Xeon* and *Itanium 2* microprocessors and AMD's *Opteron*, as well as IBM's *POWER* and *Cell* architectures. These options enable your enterprise to access literally thousands of applications designed for any of these CPUs. They provide an open architecture that's easier to manage, with a common power infrastructure that uses less electricity and generates less heat as the data center scales up and out for improved efficiency and lower TCO. Equally as important, if not more so, blade environments enable the data center to share communications resources from a single set of communications controllers across multiple heterogeneous switches.

Unfortunately, just because the architecture supports the I/O adaptors, does not imply that the system knows how to virtualize these resources across the consolidated servers and manage them easily. In a typical sever environment, the data center administrative staff must manage the Media Access Control (MAC) address for each Ethernet port and the worldwide name (WWN) for each Fibre Channel port. These addresses are hard-coded onto each server, representing multiple, repetitive tasks for the administrative staff. This process can take hours, or even days, increasing the TCO for the deployment. A blade environment really needs to be supported by an integrated communications subsystem that knows how to virtualize these services across the servers enabling the mission-critical applications to take advantage of the architecture.

Unfortunately, there is no such thing as a typical computing environment. Every enterprise, large or small, is different; some will have a "glass house" data center; many, however, also have departments and branch offices with computing responsibilities of a smaller nature, not to mention the thousands of SMEs who do not need the same scaling requirements as a Fortune 500 company.

No matter the size, the basic functionality requirements remain the same.

IBM has designed a family of bladed systems under the *BladeCenter* brand, each tailored to customer needs. (See Exhibit 1.) These chassis all support a wide variety of servers, based upon Xeon, Opteron, POWER, and Cell processors, which can be virtualized to enable sharing of compute cycles to enhance processor utilization. The data center can also configure them with additional memory, storage, and I/O controllers to support a variety of data center environments, including Ethernet, Fibre Channel, and InfiniBand. This includes third party Ethernet switches from Nortel and Cisco, Fibre Channel switches from Brocade, Cisco, Qlogic, and McData, and InfiniBand bridges and switches from Cisco and QLogic. However, with network addresses all hard-wired to the controllers, how can the data center manage the now-virtualized mission-critical environment? For IBM, the answer is simple: *BladeCenter Open Fabric Manager*.

BladeCenter with Open Fabric Manager

IBM's BladeCenter Open Fabric Manager (OFM) provides I/O virtualization to complement the server virtualization provided by applications such as *VMware*. OFM significantly reduces server deployment times, as well as management complexity, helping to reduce the data center TCO. Each BladeCenter chassis has at least one *Advanced Management Module* (*AMM*) to control all of the components in the chassis, including all of the networking blades. **A single AMM with Open Fabric Manager enables a single LAN or SAN administrator to pre-assign MAC and WWN addresses for up to 100 BC chassis, or 1,400 blade servers, before the servers are even installed.** This will enable the enterprise to reduce deployment times from days to hours, or even minutes, depending upon the number of blade servers in the domain. With OFM, the enterprise can put a new application into production as soon as the data center staff plugs the blade in. It inherits the addresses pre-assigned to that slot. Try that with a rack-mounted server, or worse, try that in a remote office or SME with no IT person to set up the addresses. Further, with the web-based interface from the AMM, the data center can control OFM functions remotely.

When configured as a separate utility or as part of *IBM Director*, **an advanced Open Fabric Manager also provides the enterprise with improved workload performance because of automatic blade failover to a spare blade in a standby blade pool.** That pool can be either within the same chassis or within any other in the 100 chassis domain. OFM automatically detects whenever a blade fails or is simply removed, selects a replacement, and configures the new blade with the addresses assigned to the failed server, powers the blade up, and continues to operate as normal. This will enable the enterprise application set to operate in a fully virtualized environment, providing the data center with the flexibility to move workloads dynamically. Open Fabric Manager is available in two configurations: a standard version integrated with the AMM to provide I/O address management, and an advanced upgrade version, which includes blade failover as a separate utility or as an extension to IBM Director. Both implementations are available year-end 2007.

OFM is an enabler for any data center deploying an IBM BladeCenter chassis. For the department, remote branch, or SME that requires a fully-functional, integrated packaging with reduced scalability in a single box, not a watered down enterprise solution, deploying a *BladeCenter S*[1] can be a godsend. It is the missing

Exhibit 1 – BladeCenter Chassis

- ***BladeCenter E*** – Helps to maximize productivity and minimize cost with an energy efficient design;
- ***BladeCenter H*** – Building on the *BC-E* design, delivers increased performance, protecting enterprise investment;
- ***BladeCenter S*** – Extends the benefits of BladeCenter design beyond the scope of the data center, into remote offices and the SME environment with right-sized configurability;
- ***BladeCenter T*** – Compliant with the more rigid NEBS requirements and rich functionality of the telecommunications industry;
- ***BladeCenter HT*** – Delivers the maximum combination of high-performance and durability required in harsh conditions for industrial, telecomm, and extreme environment deployment.

Source: IBM

[1] See **The Clipper Group Navigator** dated June 30, 2007, entitled *IBM's BladeCenter Family – The Right Choice for the SMB or Enterprise*, which is available at http://www.clipper.com/research/TCG2007070.pdf.

piece of the distributed server puzzle. IBM designed the BC-S as an integrated server for the enterprise or SME with a less demanding compute environment. With up to 6 blade servers, rather than the 14 of its big brothers, BC-S has the right computational power for a mid-sized data center environment. With up to 12 disk bays, BC-S can support 3.6TB of SAS storage or 12TB of SATA, or a mix of the two. Most importantly, with OFM, the small or mid-sized data center can effectively manage a myriad of communications combinations to meet the demands of the Internet age. BC-S with OFM also provides a conclusive nail in the IBM vs. HP comparison for the SME space.

HP recently announced the *BladeSystem c3000* as an all-in-one solution for the SME. How does BladeCenter S compare?

- **Blade Capacity** – With up to six blade server slots, BC-S compares favorably with the c3000 that supports either four full-height or eight half-height blades.
- **Storage** – BC-S can support up to 12 HDD slots dedicated to up to 9TB of storage. The c3000 uses storage blades supporting only 876GB of disk, each. These blades are mutually exclusive with HP's server blades, thus limiting the number of servers that the c3000 can support. In addition, the IBM *HS21* server blade also supports a pair of 2.5" solid-state disks that consume as little as 1 watt per drive per hour.
- **Backup** – The BC-S is ideal for typical backup methodologies. The data center can backup to high-capacity SATA disks and later move data to the supported external tape libraries, such as the *TS3100* autoloader. The c3000 must add an external library for backup support.
- **KVM Support** – The BC-S is configured standard with integrated KVM ports within the AMM. The HP c3000 is not.
- **Networking** – As previously discussed, all BladeCenters, including BladeCenter S, come with a complete set of open systems communications options. The c3000 with HP's *Virtual Connect* architecture lacks scalability and only works with their proprietary *Virtual Connect Ethernet Module* or the *Virtual Connect FC Module*. HP only manages a single chassis with each sign-on.
- **Environmental** – Acoustically, the BC-S is significantly quieter than the c3000. This is because BladeCenter S uses new fan-pack options, specifically designed to run more quietly in a small office environment. The c3000 uses the same fans as the c7000 enterprise model, accounting for the higher noise levels as compared to BladeCenter S.
- **Cost** – The BC-S costs slightly less than a c3000. As the data center adds additional servers, storage and I/O, the c3000 is twice the price of the BC-S in a more fully configured deployment.

Conclusion

Today, all data centers are looking for the same characteristics when considering the acquisition of a new I.T. solution. It must be low cost with high-performance; compatible with the broadest offering of industry-standard components to make it easy to deploy and manage; and it has to be environmentally sound, in both power consumption and its ability to coexist in a human-friendly atmosphere. BladeCenter with Open Fabric Manager fits this specification to a "T". It is open and supported across all BladeCenter chassis, with all BladeCenter interconnects.

With its cable-less design, BladeCenter, eliminates complexity and reduces deployment time. With Open Fabric Manager, BladeCenter enables the virtualization of your environment, enabling the data center to move workloads dynamically across the enterprise. It simplifies networking while interconnecting up to 1400 blades into a virtual express train.

No matter which BladeCenter chassis you select, Open Fabric Manager enables you to get on-board easily. BladeCenter S, however, holds particular interest for the SME or remote enterprise office. With its broad availability of communications engines, BladeCenter with OFM can make the connection, no matter how nodes you need to communicate to, or times the enterprise needs to switch direction. **If you need to deploy an enterprise-level solution at the right price, BladeCenter with Open Fabric Manager may be just your ticket to stay on track.**

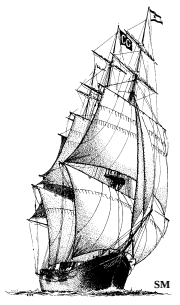

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

David Reine **is Director, Enterprise Systems for The Clipper Group.** Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial "123" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, **The Clipper Group Explorer**, **The Clipper Group Observer**, **The Clipper Group Captain's Log**, **The Clipper Group Voyager**, **Clipper Notes**, and *"clipper.com"* are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *"Navigating Information Technology Horizons"*, and *"teraproductivity"* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.