THE CLIPPER GROUP
Navigator™

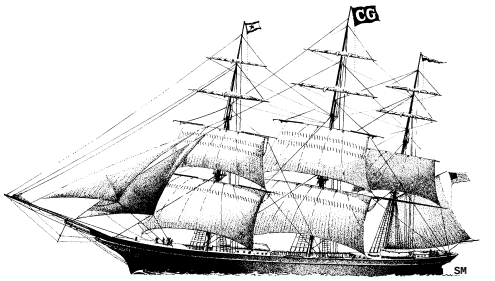

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2005015 March 22, 2005

Egenera — Ready for Prime Time in the Enterprise Data Center

Analyst: David Reine

Management Summary

Choice. We all look for it. We look for options at the ballot box when we elect our politicians: democrat, republican, or independent. We know we can only vote for one. We look for options when we select our next automobile: American or Foreign, General Motors or Toyota. We might want to choose an economical vehicle for commuting to work, a roomy one for family trips, and one more, a sporty one, to sate our mid-life crises. However, we know that we can afford only one, so we look for the strategic compromise. There are times, however, when we have a lot of choices to make, and we can have more than one. Vanilla, chocolate, and butter pecan are all delicious ice cream flavors. We can choose just our favorite or we can buy them all a pint of this, a pint of that. The cost is reasonable; the quantity is enough to share. **What would be great is having an *endless supply* of whatever flavor we felt like having, *right now*.**

Every CIO of every enterprise faces a similar dilemma in selecting the kind of information technology (think servers) and quantities needed to satisfy the mission-critical applications in the data center for various departments and organizations. There is a multitude of x86 solutions from which to choose. You have platforms based upon the 32/64-bit *Xeon* and 64-bit *Itanium* processors from Intel and platforms with 32/64-bit *Opteron* processors from AMD. You have *Windows*, *Linux*, and *Solaris* all running on x86 servers. You even have a number of flavors of *UNIX*. With virtualization, the data center can generate multiple iterations of the same operating system and/or a heterogeneous mix of different operating systems on a single server. Unfortunately, in some cases, the CIO must choose a single architecture to satisfy the demands of both CPU intensive and I/O intensive applications. In scale-up solutions, such as the HP *Superdome* or the Unisys *ES7000*, a single microprocessor, Itanium, is chosen. In scale-out environments, currently populated with 2-way and 4-way servers, you can choose from a variety of microprocessors but you must accept duplicated, and expensive, infrastructures in each node. **Blade solutions provide the best of both worlds: the ability to choose from multiple processor architectures while economizing with a common infrastructure to share between them**. Blade solutions are available from a wide variety of vendors, including all of the usual suspects: HP, IBM, Sun, and Dell.

Then there is Egenera, a company with a unique blade architecture, BladeFrame, designed to reduce data center complexity. However, Egenera believes that there is more to blade computing for the enterprise data center than just hardware. To learn more about Egenera and data center virtualization via their Egenera *Processing Area Network* (*PAN*) architecture, please read on.

IN THIS ISSUE

Utility Computing in the Data Center

Today, the enterprise data center is faced with the same requirement that is prevalent throughout the organization: *doing more with less*. For the CIO, this means reducing the total cost of ownership (TCO) for the data center through the improvement of server provisioning and server utilization. Currently, most data centers provision for peak loads in order to satisfy user demand at its highest level, for end of month or seasonal application access. By definition, high-demand periods are rarely experienced. Under normal operating procedures, each server, in fact, each processor, averages around 35% utilization, or less, to allow for the huge spike in processing that might just happen, someday. Most IA86 architectures in the data center do not have the ability to share a common infrastructure in order to increase the flexibility to respond to changing business needs in an on-demand, or "utility" environment, at least not in an easy to use, economical way.

We have all heard the term *utility computing*, but what does it mean?[1] Every CIO would love to have a data center that always had enough resources, processing, storage, and network, to satisfy the demands of every user for every application, all of the time. This would be the same environment that we have in everyday life, when we turn on a light or pick up the telephone. There is always ample electricity to meet our needs; there is always a dial tone to enable a call. The Electric Company and the Telephone Company are called *utilities*. In the data center, we can strive to have a utility model, i.e., the ability to assign resources to the hot spot, dynamically, to satisfy the service agreements that the enterprise customers expect – the utility data center.

Today's data center cannot meet this goal because of the complex nature of the system architecture currently in use, with multiple IA86 server drawers in a rack, both Intel and AMD, each with their own storage and infrastructure. In fact, complexity is the #1 problem for the enterprise data center, resulting in:

- High costs for expensive RISC servers, such as SPARC- and Power-based computers, needed to support the high availability (HA) and disaster recovery (DR) requirements of mission-critical applications;
- High capital and operational costs due to over-provisioning and underutilized servers;
- Unattainable service levels; and
- Lengthy times for application deployment, contributing to further strains on the budget.

Exhibit 1 - Data Center Requirements

- Reduce complexity for mission-critical application set and simplify operations through consolidation;
- Improve the availability of the data center compute environment;
- Optimize system performance and reduce costs through improved server utilization;
- Improve server provisioning aand reduce application time to market through on-demand system resources;
- Improve the quality of service (QoS) to the user community; and
- Lower the total cost of ownership (TCO) of network resources.

Reducing complexity and improving the ease of use of IT resources will lower the TCO to the data center and improve the enterprise's competetive advantage. But how? One way to reduce complexity is through consolidation and a shift to open source[2] applications. For many enterprises, reducing server count and implementing Linux solutions can reduce the complexity of the data center architecture, lowering the TCO by reducing maintenance charges, software-licensing fees, and administrative overhead. However, an enterprise-class solution supporting mission-critical applications, such as OLTP and ERP, requires a redundant infrastructure to deliver the needed mission-critical environment. Migrating to commodity hardware can deliver some desired TCO benefits, by reducing the cost of processors and memory. However, the mission-critical application still requires high availability and

[1] See **The Clipper Group Explorer** dated June 20, 2004, entitled *Stepwise to Utility Computing – A Pragmatist's Approach* and available at http://www.clipper.com/research/TCG2004053.pdf.

[2] An application set in which the source code is available to the general public for use and/or modification from its original design free of charge and will run on all variants of the IA86 architecture to provide access to all functionality and data through open-component based APIs and an XML-based file format. These applications are usually run under control of Windows on Intel (referred to as Wintel) or Linux on Intel.

Copyright © 2005 by The Clipper Group, Inc. Reproduction prohibited without advance written permission. All rights reserved.

reliability features that do not come standard in your average Wintel server farm. The data center is looking to lower TCO, but not to eliminate 7x24 access or the high performance that mission-critical applications require. Previously, a high-end UNIX system was required to meet these goals, which was often both proprietary and expensive. These now can be met with a commodity solution for those enterprises that have standardized on IA86 and Linux. (See Exhibit 1, page 2) **The data center staff is looking for a commodity-based mainframe-like solution to satisfy their processing and reliability requirements in a utility environment**.

Today, however, even within a commodity framework, there are different CPU architectures to choose from, depending upon the application requirements. A 32-bit architecture may suffice for some legacy applications, while a 64-bit environment may be needed for addressing a large database. Intel and AMD both offer 32-bit and 64-bit solutions, *Xeon* and *Itanium* from Intel, and *Athlon* and *Opteron* from AMD. No single processor environment will suffice for most enterprise data centers. The data center requires an integrated server capable of consolidating multiple solutions within a scale-out framework in order to maximize the enterprise resources.

To achieve this nirvana, many enterprises are migrating off of their *UNIX* platforms. Many early adopters are confirming Linux as a viable open source solution for mission-critical applications, not just infrastructure services such as file and print. Others are looking at *Windows Server 2003* and *Solaris*, Sun's UNIX implementation, in order to migrate off of existing proprietary solutions and simplify their data processing operations. Moreover, many enterprises prefer to implement their new solutions on a blade architecture, sharing resources between multiple servers over a high-speed, low-latency interconnect. Egenera has combined these requirements into an integrated solution for the data center.

Who is Egenera?

Egenera was founded in March 2000. From Day 1, they focused on solving the problems of the enterprise data center, the training ground of founder and CTO Vern Brownell. Mr. Brownell was CTO of Goldman Sachs during the 1990s where he and his staff of 1,300 were responsible for worldwide technology infrastructure, including data centers, networking, telecommunications, and trading-floor operations. He developed a deep appreciation of how the complexity of legacy server deployments in the data center and throughout the network increased costs and compromised agility. He became convinced that nothing short of a totally new approach to computing could solve the full range of problems he repeatedly encountered. He devised the *BladeFrame* system and its *Processing Area Network* (*PAN*) architecture, and formed Egenera to transform his concepts into reality.

Egenera began shipments in early 2002, signing 24 large customers by the end of 2003. Egenera looks to early adopters for acceptance of BladeFrame, the same enterprises that are leading their data center's migration to Linux. Their customer base grew to 80 worldwide by the end of 2004, with approximately 5000 servers installed. Egenera has 270 employees, with 40 in support roles. They have an agreement in place with Unisys to provide field support staff worldwide. Their customer list includes major players in a cross section of industry segments, including Financial (Commerzbank), ISP Hosting (AOL), and Government (Dep't. of Homeland Security).

In February 2005, Egenera expanded the capability of BladeFrame to include 2-way and 4-way Opteron blades, including the latest AMD64 microprocessor set, the Model 252 and 852, with the capability for an 8-way node. This was in response to customers who had already installed BladeFrame with Intel processors and were looking to take advantage of the flexibility and performance provided through Opteron. They were looking to provide their base with choice, expected from a technology leader in utility computing.

Egenera BladeFrame Solution

Egenera realized that heterogeneous commodity processors such as Xeon, Itanium, and Opteron, could be used to create commodity blades to enable modularity, but they do not solve the architectural problems that result from the complexity caused by integrating third party management software for cluster control. Egenera implemented an integrated solution via a new server architecture in the BladeFrame backplane, along with 2-way and 4-way scale-out blades to provide an ideal consolidation platform with a single management console, for simplicity and flexibility. The BladeFrame thus enables the data center to leverage Linux and Windows for mission-critical applications today, looking to

Copyright © 2005 by The Clipper Group, Inc. Reproduction prohibited without advance written permission. All rights reserved.

support Solaris by mid-2005 to satisfy customer requirements to migrate off of proprietary Sun UNIX platforms.

Egenera provides BladeFrame in two configurations: a compact *ES* implementation with six processor blades, or *pBlades*, and a full enterprise version with up to 24 pBlades. BladeFrame eliminates the expensive start-up costs of complex scale-up architectures with a relatively inexpensive secure, virtualized computing environment. The PAN architecture assembles key computing resources, processing, storage, and networking, into integrated resource pools, which can be siphoned off in support of compute services, dynamically, according to predetermined rules, and then returned to the pool when the service is completed. PAN Manager software virtualizes critical datacenter components into an integrated computing fabric. During every stage of the application lifecycle, utility computing enables the data center to lower their TCO through the reusability of adaptable resources. BladeFrame can automatically reassign resources based upon enterprise priorities and Quality-of-Service agreements and provides event-notification on a 7x24 basis. BladeFrames come with a one-year warranty for parts and labor, and telephone access to technical support, for both hardware and software issues, on a 7x24 basis.

BladeFrame is built around a very high-speed *BladePlane* at 2.5 Gigabits per second (Gbps) over redundant, serial buses to provide high-performance, low-latency communication over a secure, point-to-point fabric. A configuration consists of three types of blade drawers: a varying number of customer-replaceable Processing Blades or *pBlades*, and mandatory blades, Control or *cBlades*, and Switch, or *sBlades*, with full redundancy for all components. It is a stateless, diskless computing system working with interchangeable processing resources. It consolidates all I/O requests for IP and FC into an integrated virtual networking environment complementing the enterprise SAN strategy. BladeFrame enables built-in high availability via the PAN[3] eliminating the need for failover servers and redundant networks, complex clustering software, and the professional services required for HA implementation. This can save the data center budget big bucks that go straight to the bottom line of the enterprise financial statement

The operating environment is another example of Egenera flexibility, using Red Hat's *Enterprise Linux AS* and Novell's *SUSE Linux Enterprise Server*, and Microsoft's *Windows Server 2003*, both *Enterprise* and *Standard Editions*. Egenera developed the system management capabilities provided by PAN Manager.

From a technical standpoint, BladeFrame helps increase server utilization and eliminates the need for over-provisioning to handle peak loads or HA. **From a business standpoint, BladeFrame lowers TCO through utility computing and increases enterprise agility to respond to changing business needs.**

The major components of the BladeFrame are as follows.

pBlade

Each dual- or quad-CPU Processing Blade contains only CPUs and memory in a 1U format with no external connections, eliminating components with higher failure rates, increasing the MTBF. This architecture protects the enterprise investment in BladeFrame by enabling the easy upgrade to higher speed microprocessors. pBlades are currently available for Intel's *Xeon* and AMD's *Opteron* processors, with *Itanium* support planned for mid-2005, and are interchangeable and customer replaceable. They enable automated allocation, repurposing, and failover, with any pBlade capable of assuming the identity of any other.

cBlade

The Control Blades host PAN Manager and are redundant for high availability in an active-active state, enabling the high performance and reliability needed by mission-critical applications. They manage BladeFrame performance and functionality. Each cBlade contains multiple high-speed interfaces for 10/100/1000 BASE-T and Gigabit Ethernet connections, as well as multiple 2Gb Gigabit Fibre Channel connections. They share external traffic, enabling the pBlades to use the resources of both, simultaneously. Each cBlade also contains a DVD drive for management software portability.

sBlade

The redundant Switch Blades provide high performance point to point switching of internal network traffic for the pBlades

[3] The PAN Manager provides built-in monitoring of both the hardware and application set. Integrated management software automatically detects failures, selects the appropriate backup resources, restarts the applications and remaps the connections.

Copyright © 2005 by The Clipper Group, Inc. Reproduction prohibited without advance written permission. All rights reserved.

Exhibit 2 –
PAN Manager - Virtualization, Services and Management

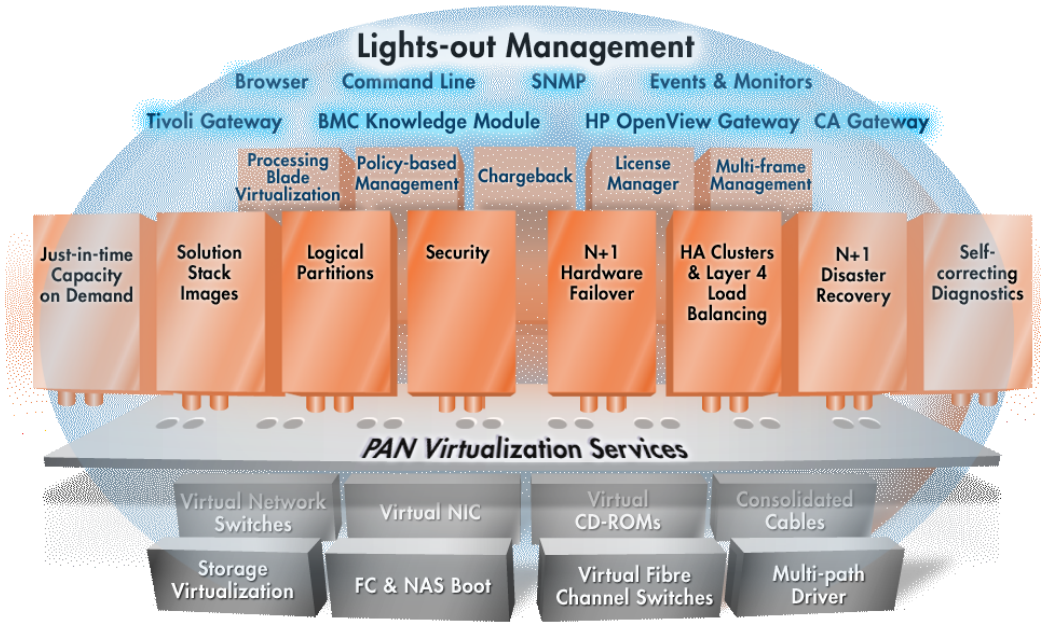

Source: Egenera

Egenera Processing Area Network Architecture

The PAN architecture (see Exhibit 2, above) provides the comprehensive system management functionality, including lights-out access and user-friendly console interface, to virtualize the data center environment. It replaces physical resources, blade and non-blade, with software equivalents – virtualizing 80% of the I/O devices that would have to exist with a legacy architecture, allocating them dynamically as enterprise demands change, according to pre-determined business service policies. PAN Manager monitors the health and performance of all system resources, enabling the BladeFrame to be self-tuning and self-healing. It manages a pool of up to 96 enterprise-class x86 processors, allocating them dynamically.

Conclusion

By combining the scalability of the BladeFrame, the price/performance advantage of x86 processors, and mainframe-class functionality, Egenera can deliver enterprise performance and reliability using commodity products such as Linux, Windows and, soon, Solaris for mission-critical applications.

Egenera's BladeFrame reduces the total implementation time of any new application, optimizing the access to system resources throughout the application lifecycle, ensuring that sufficient processing power and memory are available to the application, whenever required. The high-availability functionality of BladeFrame ensures 7x24 access for mission-critical applications. All of this lowers the TCO for the data center.

Egenera is delivering adaptable, high-quality computing for the enterprise at a lower cost. The CPU is the data center's choice, not Egenera's. The operating environment is the data center's choice, not Egenera's. The amount of money that the data center wants to save is also up to the data center. Egenera has led the enterprise to the Promised Land; it is up to you to consider the next step.

Copyright © 2005 by The Clipper Group, Inc. Reproduction prohibited without advance written permission. All rights reserved.

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**, 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.

Copyright © 2005 by The Clipper Group, Inc. Reproduction prohibited without advance written permission. All rights reserved.