


THE CLIPPER GROUP Navigator™



Published Since 1993

Report #TCG2010036

July 22, 2010

Delivering IT's Server Cloud — Fujitsu's PRIMERGY CX1000 Serves as Scalable Infrastructure and Orchestrator

Analyst: David Reine

Management Summary

Have you ever stopped to figure out exactly how many electronic components are entertaining you in your family room? From the cable modem, which is your gateway to the digital world of the Internet, to your home P.C., to your home entertainment center, there are hundreds or even thousands of small electronic parts just waiting to fail at the most inopportune time. If every failure required a visit to your home by a service technician to repair the failed component, you could be waiting for days to see him and then many hours while he tried to diagnose and repair the problem on-site. Today's devices, however, use a swap-out philosophy: you send the failed component back to the factory and a replacement can be dispatched overnight, returning you to an on-line mode as quickly as possible.

This same philosophy can also be adapted to the enterprise data center. With hundreds, perhaps even thousands, of servers deployed throughout the enterprise, swapping out a failed server with a replacement makes a tremendous amount of sense. However, there are certain implications to that plan. The data center needs to have applications that can migrate from one node to another, ideally using a virtualization engine to migrate them. There need to be multiple nodes installed with the resources available to support additional applications dynamically. With I.T. budgets already stretched to the limit, these nodes must be inexpensive, with a low total cost of ownership (TCO). This implies a dense solution to minimize floor space and an energy efficient solution to minimize power consumption.

The data center is currently experiencing an evolution of the infrastructure architecture. Today's I.T. environment requires a synergistic portfolio of platforms, applications, and services which can take advantage of consolidation, virtualization, and a cloud infrastructure. Some enterprises are looking to reduce TCO by outsourcing their peak load requirements, using the cloud to satisfy spikes in user demand, trying to maintain fiscal control above a consistent level, looking for limitless scale-out capacity beyond the financial capabilities of all but the largest enterprises. Others are looking to the cloud to provide an alternative compute capability for their high performance computing requirements. One solution that can be used in either scenario has been developed by Fujitsu to deliver a simplified Infrastructure as a Service (IaaS), with efficient access, to the cloud.

Using innovation that has evolved from decades of server development, Fujitsu has delivered the *PRIMERGY CX1000 System*, along with their *ServerView Resource Orchestrator*, for the cloud computing environment. The PRIMERGY CX1000 lowers acquisition expense and also lowers the TCO for the operation of the cloud. In the battle of price versus performance, Fujitsu asks: *why not both?* By eliminating configuration redundancy, sharing a common power and cooling infrastructure, and innovating with a unique cooling architecture, they rewrite the rules for data center economics and data center design. To learn more about the PRIMERGY CX1000, please read on.

IN THIS ISSUE

➤ Easing Burden on Enterprise Data Center.....	2
➤ Infrastructure as a Service.....	2
➤ Fujitsu PRIMERGY CX1000.....	3
➤ Conclusion	5

The Clipper Group, Inc. - Technology Acquisition Consultants ♦ Internet Publisher

One Forest Green Road ♦ Rye, New Hampshire 03870 ♦ U.S.A. ♦ 781-235-0085 ♦ 781-235-5454 FAX

Visit Clipper at www.clipper.com ♦ Send comments to editor@clipper.com

Easing Burden on Enterprise Data Center

Server sprawl is having a serious impact on the complexity of data center operations. With the IT staff trying to increase performance by deploying a single application per server, they have over-provisioned the data center with under-utilized platforms, wasting more than 80% of server compute capability and 80% of the energy required to run and cool the data center. Not only have these servers increased technical complexity, they have added significantly to the TCO of the data center. In order to regain control, the IT staff has initiated any number of programs to consolidate and virtualize data center operations to run the daily business- and mission-critical applications.

Unfortunately, the data center has multiple spikes in user demand over the course of time, still requiring the IT staff to over-provision in order to be able to satisfy every SLA agreed to. How can the data center establish a base-level of production within the glass house and still have sufficient processing power to handle incremental demand. The answer appears to lie in access to an external compute cloud, public or private, for additional application performance requirements, enabling the data center to maintain a reduced asset base. Alternatively, the enterprise could use an external cloud for all of its compute requirements, eliminating the necessity of maintaining an in-house data center. Cloud providers require high performance and unlimited scalability, along with a lower TCO from power, cooling, floor space, and administrative overhead. Cloud providers must also mask the complexity of an on-demand environment from application administrators, eliminating a requirement to deal with IT infrastructure while providing scalability to assured service levels.

Consolidation and virtualization on standard, high-performance, low-cost servers enable both the data center and cloud providers to make better utilization of data center resources. This in turn leads to a need for workload optimization, to be able to migrate applications dynamically from one server to another in order to balance workloads and enable applications to recover from system outages dynamically. A mix of physical resources in the enterprise data center and virtual resources from the cloud enables the data center to achieve a better level of cost control over variable costs for incremental workload expansion. These variables include floor space, energy con-

sumption, and administration and management. In order to control the TCO, the data center needs to have access to infrastructure services, hardware and software as well as administrative services that are external to the glass house.

Infrastructure as a Service

Cloud computing is many things to many people, in the same way that people see different shapes and images in atmospheric clouds, depending upon their point of view and imagination (and the passage of time). Some data centers view the virtual cloud as a means of controlling costs by outsourcing application processing; others see the cloud as a dynamic vehicle for accessing *additional* processing capability (in support of periodic spikes in processing requirements). Regardless, cloud computing represents a paradigm shift for providing application and storage resources to the enterprise, i.e., by providing *infrastructure as a service (IaaS)*. IaaS serves as an alternative for the data center, providing more compute cycles than are available from the existing physical infrastructure – on a temporary basis. IaaS enables the data center to have dynamic access to a virtually bottomless pool of compute cycles, paying a fixed price on a consumption basis¹.

Clouds are available on a public or private basis for a variety of uses. In addition to providing variable application demands for a user community that extends to the Internet, IaaS can support the development of new applications to meet new business requirements. It can also be used for testing and business continuity functionality to meet time constrained deadlines for processing cycles. IaaS also enables the enterprise to maintain a secure enterprise data center while meeting the mobility demands of an Internet-based community. Remote users can access the cloud to satisfy application requirements without putting the secure data center at risk, all while relying upon the cloud for security.

What is needed to meet the requirements of an IaaS environment? In addition to a highly scalable server platform, the cloud will have to provide managed services applications to a broad spectrum of clients, as well as managing pools of storage in support of the applications running in

¹ See the issue of *Clipper Notes* dated April 22, 2009, entitled *Understanding the Dimensions of Cloud Infrastructure in Order to Harvest the Benefits*, and available at <http://www.clipper.com/research/TCG2009021.pdf>.

Exhibit 1 – IaaS Cloud Requirements

1. Limitless on-demand scalability;
2. Highly-automated delivery of system resources;
3. A provisioning system to provide IT resources on-demand;
4. Ability to mask underlying technologies from the user community;
5. Provision of both role and tenant-based system administration; and
6. Delivery of complete system management capabilities for the cloud.

the cloud. One company paying close attention to the needs of the cloud is Fujitsu. With the *PRIMERGY CX1000* and their *ServerView Resource Orchestrator (ROR)*, Fujitsu has developed a total platform solution that is out of the clouds and down to earth. (See Exhibit 1, above, for a list of cloud requirements.)

Fujitsu PRIMERGY CX1000

The PRIMERGY CX1000 is an innovative platform optimized for the cloud server infrastructure and for HPC environments with a focus on the massive scaling of x86 servers, in response to a continuously rising demand for aggregated data center compute capacity. It has been designed to meet the scale-out requirements of large data centers, ISPs, ASPs, and hosts serving Web 2.0, virtualization, and IaaS environments. Deployed in a 42U rack with 38 PRIMERGY CX120 server nodes, scalable to hundreds, or even thousands, of industry-standard, fan-less, x86 Intel Xeon servers, the PRIMERGY CX1000 lowers the TCO of the IT infrastructure with an innovative *Cool-Central* architecture, which can reduce power/heat energy costs by up to 20% (and possibly more). The PRIMERGY CX1000 also includes three 2U vertical bays for a maximum of five switches, with Brocade the current qualified offering (which is discussed on the next page). Costs are also lowered with fully factory-assembled units to reduce deployment time. The starting cost of a complete PRIMERGY CX1000 is \$90K for an enclosure with 38 systems.


Cool-Central is a cooling infrastructure that enables the data center to change the paradigm of data center design, optimizing server density, reducing power consumption, and eliminating

heat dissipation problems. (See Exhibit 2, at the right.) With two central exhaust fans, the PRIMERGY CX1000 uses thermo-dynamics to purge heat through the top of the rack, where it is picked up by the data center's air circulation system. This eliminates the necessity for hundreds of local server fans and significantly reducing acquisition, power consumption and associated heat dissipation costs. This also eliminates the need for potentially costly liquid-cooling technologies. Cool-Central enables the data center to deploy PRIMERGY CX1000 frames "back-to-back", eliminating the necessity for a "Hot Aisle" (for easier removal of heat) and saving up to 40% of the data center floor space, compared to traditional rack installations. Cool-Central is part of Fujitsu's Green Policy initiative, designed to help businesses become more compliant to environmental concerns. In addition, back-to-back installations enable the IT staff to simplify maintenance operations with all front-access servers and no door.

PRIMERGY CX120 S1

Each PRIMERGY CX120 S1 node is a dual-socket server supporting a variety of Intel Xeon processors, from the quad-core *x5500*, known as *Nehalem-EP*, to the quad-core or six-core *x5600*, known as *Westmere*. The PRIMERGY CX120 can be deployed with an energy-efficient con-

Exhibit 2 – CX1000 Cool-Central Air-Flow



Source: Fujitsu

figuration consisting of dual Xeon L5630 CPUs, running at 2.13GHz, consuming only 40W per CPU, up to a high-performance configuration using dual Xeon X5670s, running at 2.93GHz with an energy consumption of 95W². Fujitsu has implemented a dual socket design in order to minimize cost and complexity, eliminating redundant components from the PRIMERGY CX120, building them into the rack with a shared infrastructure, and taking advantage of virtualization to move applications to operational platforms and restart them. The data center can then replace failed nodes, returning them to the factory for repair offline, a common practice for IaaS environments.

The PRIMERGY CX120 employs a highly-efficient local power supply, providing 92% efficiency, conforming to the *80 plus GOLD PSU* standard certification³. As described above, the PRIMERGY CX120 cooling environment has been moved to the rack infrastructure with the innovative Cool-Central technology, eliminating the necessity to install redundant fans in every node. The PRIMERGY CX1000 rack has integrated AC power distribution cabling so that the PRIMERGY CX120 nodes can be connected to power by simply plugging them into the rack.

Each PRIMERGY CX1000 rack consists of 38 factory-installed, customer replaceable 1U PRIMERGY CX120 nodes, with up to 76 Intel 5500 or 5600 CPUs. Moreover, the PRIMERGY CX1000 comes with factory assembled pre-cabling for IP connections, providing colored and pre-labeled Cat 5 Ethernet cable harness schemes for up to five switches. This enables instant setup and conflict free self-services for physical network connectivity. Each high-performance node with dual X5670 processors configured provides the data center with 456 cores and 912 threads for maximum compute power. All PRIMERGY CX120 servers have up to eight DIMMs, supporting up to 64GB of advanced ECC memory and two hard disk bays for the deployment of 2.5" non-hot-plug SATA drives, up to 500GB each. Each node also has five USB 2.0 ports (four external), three 1Gb/s Ethernet ports⁴, and one serial connector. There

² The CX120 can also be configured with Xeon 5500 processors @ 2.13GHz, 2.40GHz, and 2.93GHz.

³ 92% efficient means that 92% of the energy consumed by the server is actually used by the server, reducing wasted excess heat.

⁴ Depending on configuration choice, the CX120 is pre-cabled at the factory, minimizing data center set-up.

is one PCI gen2 expansion slot available on the PRIMERGY CX120 cloud server node for other I/O options. The PRIMERGY CX120 is qualified with Microsoft *Windows Server 2008 R2* and *Red Hat Enterprise Linux*⁵.

Fujitsu offers a full line of IT services, from concept design to operation, as well as professional services tailored to the needs of massively scale-out data centers. However, Fujitsu's most significant contribution to cloud administration lies in their *ServerView Resource Orchestrator*, the software that extends the functionality of the PRIMERGY CX1000.

ServerView Resource Orchestrator

ServerView Resource Orchestrator is a powerful management utility capable of automatically provisioning flexible cloud-ready IT architectures, enabling the IT staff to deploy a new system 90% faster than with a manual configuration, on-demand. ROR enables customers to build and install cloud-like platforms, creating an infrastructure resource pool faster and more easily than before. It is targeted at enterprises that are looking to benefit from greater IT flexibility and cost savings in moving to a cloud-like environment. It has three major functionality areas:

1. Resource orchestration,
2. Resource abstraction, and
3. Resource availability

(See Exhibit 3, at the top of the next page.)

ROR allows resources to be defined at either the logical server level or with a logical platform concept, tying together a number of logical servers. It automatically assembles resources from the pool based upon a logical definition, or profile, defined by the application administrator and based on a customizable template and a pool of pre-configured system images. This "self-service" provisioning enables the IT staff to respond faster and more flexibly to changing internal support requirements, allowing the administrator to deal with resources and not technology. The global resource pool includes storage, networking, and software images, in addition to the physical servers. All of the services are role-based in order to facilitate easy and reliable administration.

⁵ Other versions of Linux are available on request.

Exhibit 3 – ROR's Key Functionality**Resources Orchestration**

- Resource Pools
- Automated Provisioning of server, storage, and network resources
- Enables fast, efficient deployment of servers

Resource Abstraction

- Logical server and platforms
- Simplified access to resources
- Role-based IT management

Resources Availability

- Automated N:1 server failover
- Simplified disaster recovery
- Enables efficient high availability for a complete server pool

Source: Fujitsu

The rapid delivery of system resources speeds up process roll-out and enables the easy and fast addition or exchange of resources into the infrastructure. Significantly, ROR enables an increased utilization of the IT environment and a more efficient use of administration personnel, leading to an improved Quality of Service. ROR has already been enabled to access and control the management layer of storage, networking, and hypervisor products from Fujitsu and its technology partners, including Brocade, Cisco, Microsoft, NetApp, and VMware.

Brocade Switch

The PRIMERGY CX1000 includes a Brocade *FCX648S* switch as part of the PRIMERGY solution. The switch is pre-tested and certified with the PRIMERGY CX1000 as the preferred switch. It is Brocade's next generation, stackable, data center class Ethernet switch, specifically designed to address the unique requirements of a cloud computing environment. It delivers a high-performance, scalable access solution in a 1U enclosure.

The FCX648S provides 48 ports with 1GbE communication support to connect servers and storage in a HPC or cloud computing environment. It also has two 10GbE uplinks to the data center network. The typical deployment for the FCX648S is within an enterprise data center where it can be used as a top-of-rack or access-layer switch.

Conclusion

Fujitsu's PRIMERGY CX1000 has been designed to address the compute needs of the cloud environment, whether public or private, and also the scalable computing requirements of an HPC data center. The economics of the cloud are different than those of other controlled processing environments. With the TCO for the IT infrastructure included, the PRIMERGY CX-1000 provides the right balance of acquisition cost and operation cost in order to service the variable compute demands being placed upon it. If your data center is establishing a cloud environment, the Fujitsu PRIMERGY CX-1000 may be the foundation that you need.



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 a Senior Contributing Analyst for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, 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.

Disclosures

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.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

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.