Clipper Notes™
VENDOR-AGNOSTIC EXPLANATIONS AND ADVICE FOR THE INFORMATION TECHNOLOGY BUYER
Navigating Information Technology Horizons℠
New in 2007 Report #TCG2007049 March 31, 2007

Reducing Cost and Improving Performance — Consolidating the Smaller Data Center

Analyst: David Reine

Management Summary

The largest, publicly-traded enterprises can lure in the most desirable executives with benefits that are not available at most of the businesses that dot the corporate landscape. Stock options are a prime example. As we have seen in recent headlines, the award of thousands of shares of stock has skewed compensation at the CxO level wildly, as executives leave with a platinum handshake and into a private jet, to whisk them around the country from one board meeting to another. That is, if they're lucky and not heading off to the less luxurious confines provided by an institution managed by the U.S. penal system. These same enterprises also make lesser benefits available to all employees. These could include healthcare, retirement, and social benefits that smaller enterprises could not possibly hope to match, given their restricted revenue streams.

The differences between the largest enterprises and small and medium businesses have also been readily apparent in the data center. A top 500 enterprise with the largest IT budget somehow always seems to receive the most attention. It has a dedicated salesperson, dedicated sales support, and, perhaps, a dedicated service rep. The service rep takes up residence in the spacious data center, to be there whenever the mission-critical applications running on the mainframe, or one of the thousands of open systems servers, fails to perform up to the service level agreement (SLA) that was part of the acquisition contract. A smaller business would not receive the same treatment. Now, however, the differences between categories are beginning to blur, as the major vendors not only assign dedicated sales teams, or a dedicated indirect channel, to the smaller customers in their territory, but they also propose solutions tailored to addressing the biggest pains in an information intensive data center. What are these pains?

The IT staff of a smaller business has the same issues with infrastructure complexity as the largest enterprise. The person responsible for the IT environment continues to be concerned with server performance and reliability, as always, but now, with server energy consumption doubling in the past five years, and no visible means to control it, attention needs to be focused on the total cost of ownership (TCO) issues. These include server utilization, the generation of heat in the data center due to proliferation of infrastructure resources, and, obviously, the cost to cool that same data center. New technologies, such as multi-core processors and virtualization, and programs devoted to consolidate compute and storage systems, can improve system utilization. To see how these can help you lower your TCO, please read on.

IN THIS ISSUE

- **The Complex SMB Data Center** 2
- **An Evolving Open System Solution** 3
- **The Arrival of New Technologies** 3
- **Conclusion** 6

The Complex SMB Data Center

How much of the energy consumed in an office environment is used by information technology (IT) infrastructure? Moreover, how much of that energy is wasted by the IT infrastructure? Starting with the desktop P.C. sitting on *every* desk, to the fax machine and printer standing in the corner, to the application servers with disk and tape peripherals attached in the computer room, IT consumes over 75% of the energy in an average office. With the recent increase in the utility rates for that energy, depending upon where you work, it can cost more to run some of that equipment for one year than it did to purchase it. You do not have to be a card-carrying member of The Green Party to see the advantage of improving the utilization of the IT environment within your smaller data center[1].

Your enterprise is under attack: not by competitors, seeking out your corporate secrets; nor by terrorists, trying to disrupt the capitalist economy. It is under attack from the complexity of out-of-control server proliferation that has resulted in over-provisioned server farms running at utilization rates well under 25%. Over-provisioned platforms, with direct-attached storage devices (DAS) connected to each, were acquired to run one specific application. They take up too much valuable floor space and waste too much of the energy required to both run the platforms and, at the same time, the air conditioning necessary to cool the data center. The proliferation of open systems x86 servers also contributes mightily to the administrative complexity that burdens your data center staff. This workload forces you to increase the staff necessary to handle the tasks involved in managing a network with a multitude of server and DAS connections. All of these factors are damaging to your bottom line. Your data center staff must find a way to simplify the deployment, maintenance, and management of mission-critical and business-critical applications.

Increased traffic, from an ever-expanding user base across the Internet, is increasing the need to communicate, in parallel, with multiple servers in an exploding server architecture. The small business executive has learned to cope with the acquisition costs associated with the increased workload, but is definitely having trouble coping with the recurring costs associated with these acquisitions, i.e., annual maintenance charges, floor space, administrative personnel, and energy, to name a few. **These can be serious constraints to growth.** These budget-busters have a major affect on the total cost of ownership (TCO) of the IT infrastructure and reduce the profitability of the enterprise. Limitations in electrical power are even influencing the ability of the enterprise to respond to mission-critical demands, especially when the local public utility tells you that they cannot supply your data center with the electricity that you need to be competitive, regardless of the cost per kilowatt.

Obsolete servers, i.e., platforms more than one generation old, using single-threaded commodity processors from Intel or AMD, traditionally measure performance as a function of CPU clock speed. A 2GHz *Xeon* server is better than a 1GHz Xeon, and a 3GHz Xeon is better than a 2GHz model. Unfortunately, the higher the clock speed, the more energy is required, often upwards of 120 to 150 watts for the CPU alone, and the more heat is generated. **With the total cost of energy rising, the cost of powering and cooling even a small data center can become prohibitive**. The existing performance-centric chip design philosophy no longer matches the needs of an evolving data center requirement for reduced infrastructure costs.

Your data center staff must find a way to change the server paradigm in order to reduce the wasted resources, both electrical and IT staff. They need to find a way to minimize the number of servers and storage devices dotting the enterprise communication landscape, and find a way to better utilize those resources that remain. One method that has been readily accepted within the confines of the enterprise data center is to consolidate the physical assets of the environment, the servers, virtualizing more of the mission-critical applications on fewer, more-performant systems. They have also consolidated the rapidly growing storage environment by transitioning their DAS devices into storage area networks (SANs), connected to the server infrastructure through a Fibre Channel (FC) network, dedicated to one large, or a pool of, disk arrays. Implementing these programs enabled the enterprise data center to change from a *scale-out* environment to a *scale-in* architecture. While these efforts lowered the hardware costs to the enterprise, they did not have as significant an impact on administration costs, as FC administration is a specialized, and expensive, skill set. This did set the stage, however, for the implementation of SANs based on the iSCSI protocol set, allowing smaller data centers to implement storage consolidation using open systems standards, improving overall system performance and utilization, while reducing power consumption and administration. This issue of Clipper Notes discusses the advantages of consolidating and virtualizing the x86 server infrastructure of the smaller data center. The storage aspects of consolidation will be discussed in other Clipper Notes.

[1] A smaller data center may be a physical location within the enterprise, e.g., "The Glasshouse"; however, here we are referring to wherever data processing activity occurs.

An Evolving Open System Solution

Within the past year, the x86-microprocessor industry has seen a rapid evolution toward multi-threaded, multi-core processors, following the lead established in RISC[2] technology by companies such as Sun Microsystems, with *SPARC*, and IBM, with *POWER*. **Multi-core designs have changed the CPU paradigm, enabling each core to run slower, and cooler, drawing less energy, and requiring less air condition ing, while increasing total CPU performance.** AMD led this paradigm shift in the x86 arena with their open systems *Opteron* CPU, first into 64-bit processing and then into dual-core. Intel, recognizing a serious threat, followed their lead, evolving a 64-bit enabled *Xeon* from single- to dual-core. Now Intel is leading the charge into multi-core x86 processors with *Clovertown*, a quad-core Xeon CPU.

Each new evolution of the processor increases performance but maintains a power envelope similar to the previous generation. AMD, in fact, has just announced a new revision of their Opteron processor[3] that reduces the power per chip from 120 watts to 95 watts for the 2.8GHz implementation, a savings of 20%. In addition, AMD announced a set of high-efficiency Opterons[4] that lower the wattage from 95 to 68 watts, a savings of over 28%. Recently, Intel has also announced a 50-watt version of *Clovertown*[5].

Bottom Line: x86-based servers continually are achieving better performance and better performance/watt - in the same or smaller footprint.

Multi-core architecture has reinforced the viability of multi-threaded environments, enabling better utilization of CPU cycles and promoting infrastructure transformation. With the introduction of virtualization techniques and the consolidation of heterogeneous application sets onto a single platform, the data center can take advantage of the idle CPU time, **the spare 75%,** to reduce the wasted resources in the infrastructure. It is not uncommon to see utilization rates of over 75% today, reversing the earlier trend. As a result, fewer processors can accomplish more with the same amount, or less, expended energy. This is similar to one way that the federal government addressed the energy crisis 30 years ago – enabling right-turn-on-red at traffic lights to reduce wasted gasoline from idling automobile engines. As a by-product, improved gasoline utilization reduced automobile emissions, improving the air quality. There is a by-product in the data center, also. A reduction in the number of servers in the data center means less heat is being created; therefore, less air conditioning is required for cooling.

The scalability of the multi-core architecture has also created a bigger demand for less-expensive mono- and dual-socket servers, with fewer four-socket, or greater, platforms in demand in a scale-out environment. The typical mono- and dual-socket server is deployed as a 1U or 2U rack-mounted drawer, or as a blade, for example, providing the enterprise with web interface resources. (Blades definitely increase the density per square foot in the data center; however, that density also increases the energy requirement for a given space envelope. The ability to pack more computing power within a given space may lead to an increase in datacenter energy requirements.)

As the performance and scalability of these engines go up, the complexity of managing sys tem networks goes down. We also see a simplified infrastructure reducing the cost of deploying a high-performance computing (HPC) environment. These advantages outweigh the concern of some executives over the additional expenditures for acquisition. **With significantly lower recurring costs, the smaller data center can lower the TCO of the IT environment and achieve an improved ROI on new investment.**

There will still be a demand for scale-up platforms in the mid-size arena, typically running a variation of UNIX in a RISC environment, such as HP's *HP-UX* or IBM's *AIX*. This is seen in mission-critical OLTP environments, running ERP or database applications, although even here the IT infrastructure is moving to a scale-out, scale-in environment, with Linux applications leading the way.

The Arrival of New Technologies

Multi-core processors, with multiple threads running on each core, are enabling the data center to consolidate multiple servers running the same application on a single platform. In addition, they also enable multiple heterogeneous applications to run not only on each processor, but also on each core. This means that you can run multiple *Windows* applications or multiple *Linux* applications on a single core. In fact, you can run both – Windows and Linux – on a single core, minimizing the number of discreet servers required in the data center. **This industry-wide trend to consolidation may have started in the enterprise data center, but now it provides the justification for investing in virtualization to help reduce the TCO of the smaller data center, also.**

---

2 Reduced Instruction Set Computing.
3 The Opteron 8220 and 2220.
4 The Opteron 8218 and 2218.
5 No availability date has been announced.

Exhibit 1 – *1-Socket* Performance Comparison

Exhibit 1	Number of Cores					
	1 core		2 core		4 cores	
Server	TPM*	$/TPM*	TPM*	$/TPM*	TPM*	$/TPM*
Dell PE2800	28,244	$1.29	38,622	$0.99		
Dell PE2900			65,833	$0.98	69,564	$0.96
HP DL380G5					138,979	$2.12

*Note: **TPM** are transactions per minute on the TPC-C benchmark published by Transaction Processing Performance Council. **$/TPM** is based upon the total cost of hardware, software, and services used in the benchmark test by each vendor. See Exhibit 3 at end of this report.*

Consolidation

Consolidation can take many forms. When considering a typical smaller data center, we look at a mid-sized enterprise with several, or many, offices: around a state, across the country, and around the world. **The first phase of consolidation may involve the redeployment of servers from a distributed architecture to a centralized one, with all of the servers retaining their roles, however, in one data center with a single, simplified administrative infrastructure.** In this more centralized environment, the IT department can gain better cost control of the resources for which they already are responsible, something that has benefited the mainframe for decades.

The second stage of server consolidation involves *improving the utilization* of the server resources by merging servers that are running the same application onto one platform, conserving data center resources, from energy to floor space. In addition, by retiring older, obsolete servers, the IT staff can gain expense relief by trading costly maintenance contracts for new warranties.

Newer platforms provide the small data center with significantly more performance within the same power envelope of a single older server, and it can replace several. In fact, we can see significant performance gains within both transactional and HPC environments. Using the Dell *PowerEdge* (*PE*) family as a reference, we can see in Exhibit 1, above, that a single-socket, mono-core *PE2800* with a 3.6GHz Xeon processor with two threads, has a transactional throughput of 28,244 TPM[6], and is often measured at under 25% utilization. A single-socket, dual-core PE2800 at 2.8GHz, with two threads per core, has a throughput of 38,622 TPM, over one-third faster, and similar under-utilization. A single-socket, dual-core *PE2900* running at 3.0GHz has been rated with a throughput of 65,833 TPM, well over twice the throughput of the mono-core implementation. **It is also important to note that the dual-core systems have a price/performance rating of less then $1/TPM, about 25% less than the mono-core**.

Dell maintains this trend with their mono-processor, quad-core PowerEdge system, improving their price/performance ratio to $0.96, on a throughput of 69,455 TPM. HP has taken a different tack with quad-core. They have beefed up their single-socket *ProLiant DL380G5* to a throughput of 138,979 TPM, twice that of the PE2900. However, with a cost of $2.12/TPM, it is more than double that of the Dell quad-core platform[7]. This is not intended to compare Dell versus HP, but to show the reader the flexibility of a quad-core Intel Xeon (Clovertown) platform in a consolidation environment, for both low cost and high performance, depending upon data center optimization objectives.

A dual-processor, single-core version of the PE2800 has a throughput of 63,646 TPM, comparable to the dual-core PE2900, but with a price/performance cost of $2.28[8], it is more than twice that of the newer engines[9]. (See Exhibit 2, at the top of the next page.) It also happens to be the same in terms of price/performance as the *dual-processor* HP DL385G2, a dual-core *Opteron* platform with a performance rating of 139,693. Quad-core systems from HP also have significantly higher performance ratings, with cost ratios two to three times that of the mono-CPU, dual core servers. Because of the poor utilization of older platforms and the enhanced performance of the newer multi-core processors, the data center can combine several older platforms into a single system.

There is only one published result for a quad-processor, dual-core implementation. However, while it has a superb performance rating in excess of

[6] Transactions per minute are based upon results from the Transaction Processing Performance Council's TPC-C benchmarks.

[7] The Dell PE2900 system as configured for this test costs $66,455, while the HP DL380G5 costs $292,542.

[8] It must be noted that the price/performance ratio does not include the cost of energy to run and cool the platform, however, it does include an optional Relational DBMS and disk storage, which add significant cost to the total configuration.

[9] Please note that each of these systems has a substantially-different configuration. The storage on each is sized to the TPM rating and often represents one-half to two-thirds of the included cost.

Exhibit 2 – *2-Socket* Performance Comparison

Exhibit 2	Number of Cores					
	1 core		2 cores		4 cores	
Server	TPM	$/TPM			TPM	$/TPM
Dell PE2800	63,646	$2.28				
HP DL385G2			139,693	$2.28		
HP BL480c					222,117	$2.72
HP ML370G5					240,737	$1.85

Note: **TPM are transactions per minute on the TPC-C benchmark published by Transaction Processing Performance Council. **$/TPM** is based upon the total cost of hardware, software, and services used in the benchmark test by each vendor. See Exhibit 3 at end of this report.*

330,000 TPM, it carries a significantly higher cost, with a price/performance rating of $5.30/TPM. It is important to understand that, even though there are eight cores in the quad-socket platform, this is not an apples-to-apples comparison with the systems discussed above (and intentionally so.) It is intended to show another configuration option, at another price point. Any vendor - configuring a quad-CPU system using dual-core processors - likely will carry a similar infrastructure burden.[10] This burden includes higher software charges for a multi-socket system.

Bottom Line: We have used a commodity benchmark to compare the price/performance of a variety of one-, two-, and four-socket platforms. As you examine the details of the vendors' configurations, you will note the impact of the socket count on the cost of both the operating environment and the data management selected. You will also note the impact that the size of the storage environment has on both the overall performance and the total cost. The higher the disk count, the greater the performance and the higher the price. Accordingly, the smaller data center may be better off deploying multiple one- or two-socket platforms than a single four-socket configuration in a transactional environment.

Data centers do not measure High-Performance Computing in transactions, but rather by computation. The Standard Performance Evaluation Corporation (SPEC) has established an extensive series of computational benchmarks to establish the relative performance of competing systems. Some of these benchmarks, such as SPEC *Cint2006* and SPEC *Cfp2006*[11], measure the single-threaded power of the microprocessor, gaining little advantage from multi-core or multi-processor architectures. Others, such as SPEC *Cint2006 Rate* and SPEC *Cfp2006_rate*, do take advantage of multi-core and multi-processor architectures and provide a fair comparison of platform performance. For example, a Dell Xeon workstation with a single dual-core CPU has a SPEC Cint2006_rate of 29.2. A quad-core *PE840* with a *Xeon* processor, however, has a rating of 43.0, 47% higher. SPEC lists comparable results for the HP *ProLiant DL585* with an AMD *Opteron 854* processor. A dual-processor DL585 has a rating of 22.3, while a quad-processor implementation comes in at 41.4, 85% higher.

The floating-point benchmarks show similar results, with a dual-processor, single-core Sun Fire X4200 (*Opteron 256*) showing a rating of 24.3 in the Cfp2006_rate test, with a dual-processor, dual-core version (*Opteron 285*) coming in at 36.0, almost 50% higher. The data center also can achieve comparable savings by consolidating storage.

Virtualization

Server virtualization is a multi-dimensional opportunity to use server resources more efficiently.[12] In the most straightforward implementation, it is a way of consolidating many older servers onto a new server, with each older server's applications running in a fixed partition. For example, a server may be partitioned into 10 virtual servers, each getting 10% of the (newer, faster) server's processing capacity. In a more complex scenario, the partitions can be dynamic, i.e., allocated only long enough for the application to complete its task.[13]

Additionally, virtualization can lower your TCO of the IT infrastructure by reducing the charges for licensing many operating systems and middleware. Microsoft, for one, licenses *Windows Server 2003* by

[10] However, if you need to do several hundred thousand TPM, or more, then you will need much more throughput than can be delivered by the lower-cost systems shown in Exhibits 1 and 2.

[11] CINT measures integer performance, while CFP measures floating-point performance.

[12] See the February 27, 2007, issue of Clipper Notes entitled *Server Virtualization Made Real*, which is available at http://www.clipper.com/research/TCG2007028.pdf.

[13] See the February 28, 2007, issue of Clipper Notes entitled *Virtual Machines - Three Things to Consider + Three Ways to Use*, which is available at http://www.clipper.com/research/TCG2007029.pdf.

the socket, not by core or replication. The data center can execute multiple applications on a single socket with one Windows license, instead of licensing Windows on multiple platforms to execute a single application.

Conclusion

As a small business executive, you have the responsibility to reduce the TCO of an open systems data center running a myriad of Windows and Linux applications. You can continue to execute them on an under-utilized server farm that has grown beyond your expectations, or desires; or you can modernize your IT infrastructure. By replacing your obsolete, and underperforming, x86 architecture with multi-core processors, you can spend a nickel to save a dime. By replacing an aging network of servers, you can eliminate onerous maintenance contracts and reduce burdensome software licenses. You can simplify the IT infrastructure by combining common applications on fewer systems, minimizing the costs to manage and to power. You can consolidate multiple heterogeneous applications that were wasting data center power, burning too much electricity and too many dollars.

The higher performance capability of the new server technology will enable you to deploy less expensive one- and two-socket servers, rather than the more costly four-socket platforms. The data center can continue to run the same application set, as the new platforms are completely compatible with your older servers, whether you go with AMD- or Intel-based systems.

As a wise man once said: “You have to spend money to make money”. In this instance, you can spend a little money to save a lot.

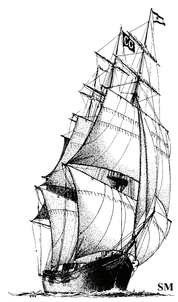

Exhibit 3 –
References Cited from Transaction Processing Performance Council, for TPC-C Benchmark
(All are copyright by Transaction Processing Performance Council.)

For Dell PowerEdge 2800 (1-Socket):
TPC Benchmark C Full Disclosure Report for Dell PowerEdge 2800 Using Microsoft SQL Server 2005 Standard x64 Edition and Microsoft Windows Server 2003, Standard x64 Edition, dated April 14, 2006.

For Dell PowerEdge 2800 (2-Socket):
TPC Benchmark C Full Disclosure Report for Dell PowerEdge 2800 Using Microsoft SQL Server 2000 Enterprise Edition and Microsoft Windows Server 2003, Enterprise Edition, dated April 14, 2006.

For Dell PowerEdge 2900:
TPC Benchmark C Full Disclosure Report for Dell PowerEdge 2900 Using Microsoft SQL Server 2005 Standard x64 Edition and Microsoft Windows Server 2003, Standard x64 Edition SP1, dated March 9, 2007.

For HP ProLiant DL380 G5/2.66GHz:
TPC Benchmark C Full Disclosure Report for HP ProLiant DL380 G5/2.66GHz Quad Core using Microsoft SQL Server 2005 Standard x64 Edition SP1 and Microsoft Windows Server 2003, Standard x64 Edition SP1, dated February 13, 2007.

For HP ProLiant DL385 G2/2.8GHz:
TPC Benchmark C Full Disclosure Report for HP ProLiant DL385 G2/2.8GHz Dual Core using Microsoft SQL Server 2005 Enterprise x64 Edition SP1 and Microsoft Windows Server 2003, Enterprise x64 Edition SP1, dated November 9, 2006.

For HP ProLiant BL480c Server Blade:
TPC Benchmark C Full Disclosure Report for HP ProLiant BL480c Server Blade using Microsoft SQL Server 2005 Enterprise x64 Edition SP1 and Microsoft Windows Server 2003, Enterprise x64 Edition SP1, dated November 14, 2006.

For HP ProLiant ML370G5 SAS/2.66GHz QC:
TPC Benchmark C Full Disclosure Report for HP ProLiant ML370G5 /2.66GHz Quad Core using Microsoft SQL Server 2005 Enterprise x64 Edition SP1 and Microsoft Windows Server 2003, Enterprise x64 Edition SP1, dated November 13 2006.

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.