Clipper Notes™
VENDOR-AGNOSTIC EXPLANATIONS AND ADVICE FOR THE INFORMATION TECHNOLOGY BUYER

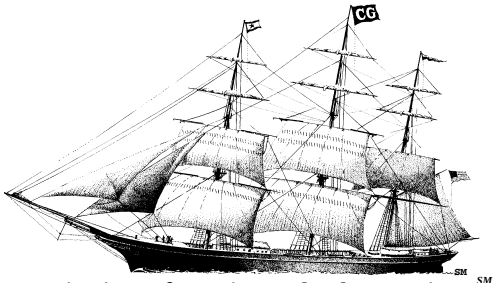

Navigating Information Technology Horizons℠

Published Since 2007 | Report #TCG2009032 | August 7, 2009

Reducing TCO of the Enterprise Data Center — Upgrading and Replacing Aging x86 Infrastructure

Analyst: David Reine

Management Summary

The average sports fan is constantly aware of the transitional nature of the roster of his/her favorite team. No matter what the sport, you can be sure that management will try to put the most competitive product on the field. To do this, the team will draft young athletes out of high school or college who may be better then the players at the bottom of their roster in order to improve the quality of the product on the field. They will also try to sign experienced veterans who are as good, or better, than players already on their team, but may be less expensive than the players being replaced, in order to address the economic nature of sports in an era of a downward spiraling economy and a fixed limit for salaries. An experienced management will also try to replace some of the aging superstars on their roster, the heart of the team, so that the entire starting lineup does not get old at the same time, removing the team from contention. The management team has to consider performance, cost, and competitiveness when fielding a professional team.

The CIO of any enterprise has to consider many of these same factors, and others, when calculating the total cost of ownership (TCO) of the IT infrastructure of the data center. For many years now, x86 servers have been driving the infrastructure services for many enterprise data centers: web, mail, file, and print to name a few. Recently, with the acceptance of Linux into the mainstream, some mission- and business-critical applications also have moved onto open x86 architectures. Unfortunately, in the same spirit of "one man, one vote", many CIOs have deployed these open systems servers with a single application on each, thus possibly employing only 15% to 20% of the CPU resources available. This means that not only up to 80% of the processing power of the server was being wasted, but also 80% of the energy required to drive the platform and 80% of the energy required to cool the data center were also being wasted, thus driving up the TCO of the IT infrastructure. When combined with the cost of floor space for each server and the software licensing fees, it is easy to see how a growing TCO can negatively impact the bottom line of the enterprise.

With the recent announcements from Intel and AMD of high-performance, energy-efficient, multi-core processors, many data centers are moving to consolidated environments, using virtualization to share processing resources between multiple applications, thus improving CPU utilization to upwards of 80%, reducing server count by up to a 10:1 ratio, or higher, and reducing the number of administrators required to manage the environment. In some cases, this consolidation has enabled an enterprise to avoid having to build new data center facilities at a cost of millions of dollars. To learn more about how your data center can reduce the TCO of IT operations with AMD's new quad-processor (*Shanghai*) or six-processor (*Istanbul*) *Opteron* CPU, or Intel's new quad-processor (*Nehalem-EP*) or six-processor (*Dunnington*) *Xeon* processor, please read on.

IN THIS ISSUE

Addressing Data Center Sprawl

In order to reduce the growing total cost of ownership (TCO) and the increased complexity of the server infrastructure in enterprise data centers, and the remote offices of all enterprises, it has become apparent that the IT staff in these data centers must consolidate the IT infrastructure and simplify daily operations, especially the deployment of new applications. In addition, due to the under-utilization of IT resources throughout the enterprise and the accompanying increase in the TCO of the IT architecture, the staff must also do whatever is necessary to optimize the IT infrastructure.

To simplify the daily operations of the data center, optimize the infrastructure, and reduce data center TCO, the IT staff must take advantage of the increased processing capability of the latest open systems CPUs, either AMD's *Opteron* or Intel's *Xeon*, to update legacy environments to eliminate the server sprawl that has overtaken enterprise resources, by consolidating the data center environment. Once the consolidated architecture has been put in place, the IT staff can virtualize multiple applications onto a single platform using improved power management, increased memory bandwidth, and enhanced virtualization capabilities, which now are rapidly becoming the trademark of x86 processors, with their larger memories, faster communication channels, and unified I/O, along with improved systems management tools. Improving the ease of use of mission- and business-critical applications, as well as daily infrastructure applications, is essential to lowering the TCO of the data center. The total cost of ownership clearly involves more than just the acquisition costs of new IT resources. It also includes the cost of energy consumption, software licensing, floor space, administration, and maintenance, all clearly factoring into any accounting equation to calculate TCO. For example, **if you were to run out of floor space in your data center, what would it cost to build another?**

Because of the backward compatibility of the x86 architecture, delivered as server CPU chips from AMD and Intel, upgrading the IT architecture to the latest revision of either Opteron or Xeon is not a complex step for the data center staff. If the IT staff needs to replace a server on a one-to-one basis in order to increase performance scalability, the upgrade is simple and straightforward. On the other hand, if the staff plans to change the paradigm of the data center by consolidating multiple applications onto a single server and virtualizing that server to make more efficient utilization of its resources, the process is a little more complex and the CIO needs to ensure that the enterprise does not have to repeat the process in the foreseeable future. Any new acquisition must have open-ended scalability to be able to handle tomorrow's challenges as well as today's. The new architecture must enable the data center to improve the performance of the IT environment while, at the same time, driving down costs. It must also facilitate the deployment of new applications and the ability of the IT staff to manage them. Implementing more cores with multiple threads is one way to do just that.

The latest round of announcements have presented the data center with four options to choose from, in terms of CPU architecture, with those CPUs embedded in a multitude of server options from all of the leading platform providers, including Dell, HP, IBM, NEC, and Sun. Depending upon the application mix, the IT staff needs to qualify the best option for their environment from among the available choices: AMD's Opteron, either the quad-core *Shanghai* or the six-core *Istanbul*[1], or Intel's Xeon, either the quad-core *Nehalem-EP*, or the six-core *Dunnington*[2]. This nomenclature refers to the manufacturer's internal code names; however, due to the possible confusion over the numerical naming conventions assigned for production by the vendors, the internal code names will be used herein for clarity.

What Can You Do?

First, do not second-guess yourself over the decisions made five years ago, or even three years ago. There were sound reasons for those decisions. Today's economic climate does not mean those decisions were wrong. **The only mistake would be to continue to live with an IT infrastructure that costs more to maintain than to replace!**

Most mid-size companies keep their servers for three to five years. In this economy, many may think that delaying a refresh is smart. But that could actually drain your already-limited IT

[1] The Shanghai version of Opteron is a *Model 2300* for dual processor systems *or 8300* for 4-to-8 CPU systems, while the Istanbul is a *2400* for dual processors and an *8400* for 4-to-8 CPU models.

[2] The Nehalem-EP is an Intel Xeon *Processor 5500 Series* for dual processor systems, while the Dunnington is an Intel *Xeon Processor 7400 Series* and scales to over four CPUs.

budget, for operational costs and other soft expenses of supporting old servers far outweigh the costs of consolidating on new equipment. Today's economic conditions suggest that your enterprise should hold on to your installed servers even longer, rather than make a significant acquisition expense (CapEx). However, doing that could actually deplete the already limited IT budget, because operational costs and other TCO expenses (OpEx, e.g., maintenance, energy consumption, software licenses, etc.) to support your infrastructure may be higher than the costs of consolidating and virtualizing your IT infrastructure on a more modern architecture. With pressures mounting to control costs, a platform refresh might be exactly want you need to restore sanity to your budget.

What needs to be done? First, and foremost the CIO must reduce the TCO of the IT infrastructure. To do that, any new server must have the performance level to consolidate and support today's application set without processing bottlenecks, and also have the scalability to grow both the application set and user count! It must be able to reduce the energy consumption level in order to reduce the demand on the data center power resources, enabling the CIO to prolong the life of the data center facilities. This requires the integrated management of power resources within the CPU.

You must be able to virtualize these platforms, and virtualize them easily, so that you do not place undue stress on systems administrators, increasing the costs in that arena. The virtualization of the IT environment is highly dependent on the ability to scale memory in order to support multiple applications. The new architecture must support sufficient memory (DIMMs) to enable the IT staff to virtualize anywhere from five-to-ten applications at a time. Furthermore, with multiple applications running, the server must have the capability to support high-performance I/O channels (8Gb F.C. and 10 GbE) in order to eliminate throughput bottlenecks[3]. Clearly, this implies support for a multi-socket environment with a high-speed communications capability between processors to facilitate a shared environment for both growth and the flexibility required by a truly virtualized architecture. In addition, the acquisition of new servers will restart the warranty clock, eliminating expensive maintenance charges from the TCO equation.

How well the processor can support these qualities will determine how far the IT staff can go to consolidate the IT infrastructure. Let's take a look at how well AMD and Intel have done in responding to the demands of the enterprise to lower the TCO of the data center.

AMD Processor Functionality

AMD's Opteron technology was introduced in April 2003 to compete with Intel's Xeon family for a share of the x86 market. Then, it was the only processor with the capability to run native execution of x86 32-bit applications as well as the native execution of x86-64 64-bit applications. AMD announced the availability of a quad-core Opteron, named *Barcelona*, with a 65nm process and new power and thermal management techniques, in September 2007. The most recent quad-core announcement introduced a processor named *Shanghai*, with a 45nm process, in November 2008. These CPUs were designed to protect data center investment in IT and enable consolidation to simplify IT management. The six-core *Istanbul* version of Opteron, also a 45nm processor, was introduced in June.

AMD Quad-Core Opteron Processor

The quad-core AMD Opteron processor was designed for not only high performance, but also for the "**highest performance per Watt" ratio** that AMD could muster, along with an integrated virtualization capability, *AMD-V Virtualization*, in support of whichever virtualization hypervisor the IT staff requires. The Shanghai processor has been designed with multiple processing points, each with its own power requirement, for up to two-way and up to eight-way servers. A **dual-socket Shanghai** is available in clock speeds ranging from 2.1 or 2.3GHz, with an average CPU Power (ACP[4]) requirement of only **40W**, in their energy efficient "*EE*" version, up to 2.4GHz with an ACP of **55W** in the "*HE*" model, up to 2.9GHz in a standard **75W** ACP package, with a high-performance version to 3.1GHz in an "*SE*" model with an ACP of **105W**. An **eight-socket Shanghai** comes in the same options, with the exception of the "EE" processors, which are not appropriate for a platform that scales to that magnitude.

Because of its integrated memory controller

[3] See the issue of Clipper Notes dated May 29, 2009, entitled *Is Your IT Infrastructure Ready for Tomorrow? – Preparing for a 10GbE Future*, and available at http://www.clipper.com/research/TCG2009025.pdf.

[4] ACP is a metric that represents real-world processor power consumption to the IT managers responsible for power budgeting.

Exhibit 1 – Opteron Features

- ***HyperTransport Technology***– Provides up to 8GB/s of bandwidth per link to connect processors and I/O subsystems, improving application performance and scalability.
- ***Direct Connect Architecture*** – Helps to reduce bottlenecks inherent in traditional front-side bus architectures.
- ***Integrated Memory Controller*** – Offers high-bandwidth, low-latency access to memory for superior performance on memory-intensive environments, as well as virtualized environments.
- ***AMD-P Power Management*** technologies – with ***Enhanced AMD PowerNow!*** technology to reduce data center energy consumption, ***Dual Dynamic Power Management*** to reduce idle power consumption and permit per-processor power management in multi-socket designs, ***AMD CoolCore*** technology to cut off power to unused transistor areas to reduce power consumption and lower heat generation, ***AMD PowerCap Manager*** to enable the I.T. staff to put a cap on the P-state level of the cores via the BIOS, and ***AMD SmartFetch*** technology to allow inactive cores to enter a "halt" state and draw less power.
- ***AMD-V technology*** – to enable the translation of virtual to physical addresses in hardware rather than software, I/O virtualization, and the migration of a virtual machine between two physical machines running AMD Opteron technology. This feature is also known as Rapid Virtualization Indexing (RVI).

Source: AMD

and *Direct Connect Architecture* (*DCA)*, Opteron improves system performance and efficiency by directly connecting the processors, memory controller, and I/O to the processor, reducing memory latency. Opteron enables simultaneous 32- and 64-bit computing, thus protecting enterprise investment in existing software. (See Exhibit 1, above, for a detailed description of the significant features of the Opteron architecture.)

The quad-core architecture is most appropriate for applications that require a balance of clock speed and threading while staying within a specific power budget in dense environments, such as web services, cloud computing, network infrastructure, and email.

AMD 6-Core Opteron Processor

AMD has just announced the availability of a six-core Opteron processor, complementing the quad-core Opteron announced last year. This CPU has been developed specifically to improve upon the features of the four-core Shanghai CPU in support of the enterprise data center's need for improved throughput, expanded virtualization capabilities, energy efficiency, and better economics resulting in a lower TCO. Designed to the same power envelope as the four-core Shanghai processor, Istanbul enables increased capabilities for an HPC environment or increased database performance, while reducing TCO by maintaining the same energy specifications as its predecessor. Where the Shanghai CPU supports a HyperTransport Technology of up to 17.6 GB/s @ 4.4 GT/s per link, Istanbul can support up to 19.2 GB/s @ 4.8 GT/s per link via *HT Assist*[5].

The innovative features in the quad-core CPU are continued in Istanbul, enabling improved performance and consolidation.

- *With six cores*, and up to 128GB of RAM Istanbul provides the ability to manage virtual machines dynamically with outstanding performance via the *AMD Virtualization* (*AMD-V*) technology that the enterprise requires to consolidate underutilized systems.
- *With six cores*, Istanbul delivers significant performance improvement over Shanghai, operating within the same power and thermal envelopes, providing outstanding performance per watt results, with no increase in data center energy consumption.
- *With six cores*, the data center can build a scalable infrastructure that will enable the I.T. staff to utilize – fully – platforms that support two, four, or eight processors, with up to 48 cores and 96 threads, for the growing needs of the enterprise.
- *With six cores*, Istanbul can replace older platforms and maintain compatibility with your existing I.T. infrastructure, running your mission- and business-critical applications with increased performance, while reducing energy requirements even further.
- *With six cores*, AMD has designed Istanbul in two distinct families: the *AMD Opteron 2400* and the *AMD Opteron 8400*. The data center can deploy the Opteron 2400, replacing the *Opteron 2300*, with a two-

[5] HT Assist reduces cache probe traffic between CPUs, supporting faster queries in 4- and 8-way servers, increasing performance for cache sensitive applications such as database, virtualization, and compute intensive applications.

socket server, while deploying the Opteron 8400, replacing the *Opteron 8300*, with a four-socket or eight-socket system. The 2400 CPUs come with frequencies of 2.2, 2.4, or 2.6GHz for low-power, standard-power, and performance-optimized requirements.

The Istanbul CPU also has been designed not only with multiple processing points, but also with multiple ACPs. A system with up to two sockets is available in five clock speeds ranging from 2.1 to 2.8GHz, while the eight-socket Istanbul comes in four options, from 2.1 to 2.8GHz, a 55W model at 2.1GHz, two standard 75W models at 2.4 and 2.6GHz, and a performance-optimized 105W version at 2.8GHz.

What does all of this mean to an enterprise trying to deploy new applications without adding to the TCO of the data center? As always, the answer is: ***It depends!*** If the IT staff is currently working with a 4-year-old environment consisting of dual-socket, single core Xeon platforms, it can mean a lot to the bottom line. One Tier-1 server manufacturer calculates that they can replace 42 two-socket platforms consuming 10,700W per year with four new platforms, each configured with two six-core Opteron CPUs[6], consuming 530W, at a cost of about $5,800 each. With a consolidation ratio of about 10:1, not only would this lower operating costs by $17,500, it would also lower software fees for the Linux operating system alone by $30,400 per year, **resulting in an ROI of approximately three months**, without considering the effects on application software, floor space, or others TCO variables. After the first three months, these savings go straight to the bottom line! If the data center was configured with 2-year-old dual-core Xeon servers, these same four Istanbul platforms could replace 14 of the older servers with an OpEx savings of $3,350 per year from power and cooling an $8,000 per year savings from software licensing. **In this case, the ROI would be 11 months, with future savings going back to the enterprise**.

Intel Processor Functionality

As with AMD, Intel has focused developments in the Xeon technology on increasing performance while reducing operating costs. While Intel has maintained their front-side bus (FSB) architecture on Dunnington, they have evolved to an integrated memory controller for Nehalem-EP in order to reduce the latency of processor commands and improve the virtualization qualities of the Xeon architecture.

Exhibit 2 — Xeon Nehalem-EP Servers

Model	Cores	Frequency	Power
X5570	Quad	2.93 GHz	95W
X5560	Quad	2.80 GHz	95W
X5550	Quad	2.66 GHz	95W
E5540	Quad	2.53 GHz	80W
E5530	Quad	2.40 GHz	80W
E5520	Quad	2.26 GHz	80W
E5506	Quad	3.20 GHz	80W
E5504	Quad	3.20 GHz	80W
E5502	Dual	3.20 GHz	80W
L5520	Quad	2.26 GHz	60W
L5506	Quad	2.13 GHz	60W

Source: Intel

Quad-Core Nehalem EP Series

Nehalem-EP provides a foundation for IT management to refresh existing or design new data centers to achieve greater performance while using less energy and space, reducing operating costs significantly. It features intelligent performance to optimize performance to fit application needs, scales energy usage to the workload, achieving optimal performance/watt, and provides flexible virtualization to improve IT infrastructure and reduce TCO.

As with the AMD Istanbul processor, the Intel Nehalem-EP processor is available in a variety of processor speeds and power levels. (See Exhibit 2, above). The functionality, however, is common for all versions. The major improvements made for Nehalem-EP include:

- **An integrated memory controller** within the processor die. This may be the feature with the biggest impact. It eliminates the requirement for a front-side bus and increases performance by lowering memory latency without increasing clock speed. This keeps the thermal envelope at the same level as the previous Xeon CPU.
- Support for up to **three channels of DDR3 memory** – increasing memory bandwidth.
- **Independent power** for each of the four-processor cores enables more control of the energy being consumed by the processor.

[6] Each server was configured with two AMD2435 processors running at 2.6GHz with 32GB of memory and a power requirement of 75W.

- **Four cores on a single die** improve core-to-core communication, with higher bandwidth.
- **QuickPath Interconnect** – a high-speed, chip-to-chip interconnect technology to enable the high performing x5500 CPUs to connect to another component or another chip on the processor board.
- **45nm process** – the smaller architecture enables the CPU to run with lower power and higher clock speed, squeezing more transistors onto the CPU die.
- **Hyper-Threading** – each core supports simultaneous multithreading (SMT), with two threads per core, for a total of eight threads, improving performance for multithreaded applications.
- **Advanced Power Management** provides support for an additional, smaller processor core, transparent to the operating system. This power management unit manages the power envelope of the 5500 via sensors built into the main cores to monitor heat, power, and current, optimizing energy consumption, as required. The Xeon 5500 also minimizes idle power, invoking a sleep state, as needed.
- **Turbo Mode** – an automatic over-clocking feature enabling individual cores to be driven dynamically at a higher clock frequency to improve application responsiveness. Turbo Mode is controlled by the power management unit and executes transparently to the operating system.

With all of these improvements, Nehalem consumes 30% less energy with the same performance as its predecessor[7]. In addition, it provides up to 2.5 times the performance for enterprise computing, 3.5 times the bandwidth for technical computing, and twice the software threads over the previous version of Xeon.

What do all of these improvements to Xeon mean to the data center in terms of TCO? As with the earlier Opteron analysis, if the IT staff needs to replace a 4-year-old environment consisting of dual-socket, single core Xeon platforms, it again means a lot to the bottom line. The same Tier-1 vendor calculates that they can replace 42 two-socket, single-core servers consuming 10,700W per year with four new systems[8], each configured with two quad-core Nehalem-EP processors, consuming 450W, at a cost of about $6,500 each. Not only would these lower operating costs by $18,000, it would also lower software-licensing fees by $31,200 per year, **resulting in an ROI of approximately three months.** After this period, these savings go straight to the bottom line! If the data center was configured with 2-year-old dual-core Xeon servers, these same four Nehalem-EP platforms could replace 14 of the older systems, at 2,630W, with an OpEx savings of $3,850 per year from power and cooling an $8,800 per year savings from software licensing. **Here, the ROI would be 12 months, with savings going back to the enterprise after that**.

Six-Core Dunnington Series

Specifically designed for virtualization, Dunnington processors provide the data center with the performance headroom required to consolidate applications onto fewer systems using proven virtualization solutions. It also provides the compute power needed for mission- and business-critical applications, such as ERP and business intelligence.

Dunnington was designed with 45nm technology, 50% more cores than previously available, *Intel VT*[9] to provide a hardware-assist to your virtualization environment, and 16MB of L3 cache per processor to help drive data-demanding enterprise applications faster via more virtual machines. Dunnington provides maximum flexibility, allowing live migration across all Intel Core micro-architecture-based servers. Dunnington provides greater performance and scalability with 4-to-32 processor enterprise servers.

As with Nehalem, the six-core Dunnington is available in a variety of processor speeds, from 2.13GHz to 2.66GHz and multiple power levels, from an energy-conserving 65W to a performance-optimized 130W. Dunnington is also available as a quad-core processor running at an energy-efficient level of 50W.

Dunnington continues Intel's legacy design of FSB processors. The Intel 7300 chipset, however, improves data movement by increasing interconnect bandwidth, optimizing system bandwidth, increasing memory capacity, and improving network traffic processing while

[7] Based on data provided by Intel.

[8] Each server was configured with two x5550 processors running at 2.66 GHz with 32GB of memory and a power requirement of 95W.

[9] Intel Virtualization Technology.

reducing I/O latency. Dunnington has enhanced reliability through a number of features, such as:

- **Memory ECC** to detect and correct single-bit errors;
- **Enhanced Memory ECC** to retry double-bit errors;
- **Memory Sparing** to enable the systems management controller to predict a failing DIMM and copy the data to a spare memory DIMM, maintaining server availability and uptime;
- **Memory CRC** to automatically retry address and command transmissions if a transient error occurs; and
- **Symmetric Access to all CPUs** to enable a system restart, if the primary processor fails.

Conclusion

First, there are a number of significant differences between Intel's Dunnington architecture and its Nehalem-EP and AMD's Opteron architecture. For example, Nehalem offers QPI[10], while AMD offers HyperTransport, both of which are direct connect NUMA[11] architectures, while Dunnington is still based on the old FSB architecture, requiring the Northbridge chipset. Additionally, for virtualization workloads, Nehalem-EP and the AMD solutions offer RVI (Rapid Virtualization Indexing) that increases performance dramatically while Dunnington does not. As such, it appears that the AMD 4-socket, 6-core designs will outperform the Intel 4-socket 6-core designs and, thus, is the primary recommendation for enterprises that are looking for leading performance and/or the lowest acquisition cost. However, Intel has already announced plans to deliver Nehalem-EX in 2010, an eight-core processor that will replace Dunnington.

In general, Nehalem-EP server may cost more to acquire but usually will provide greater operational savings versus AMD. On the other hand, the AMD-based systems may cost less to acquire but require more power to operate. Based on *SpecPower* results, the ROI payback period of three months (when consolidating single core servers) and 12 months (when consolidating dual core servers) is virtually identical between AMD- and Intel-based systems but is arrived at very differently (AMD has lower acquisition costs while Intel has lower operating costs). This being the case, we would recommend an Intel-based solution, if the enterprise is concerned about operating costs, and AMD-based solutions, if the customer is concerned about acquisition costs.

In any case, it makes no sense to continue spending budget on systems with a TCO that rises every month. Every enterprise needs to initiate a program to replace an aging, single-core or dual-core server environment immediately, to avoid throwing good money away to maintain inefficient IT architectures.

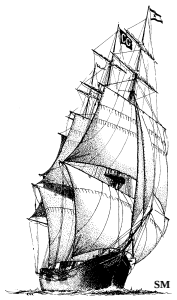

[10] QuickPath Interconnect.
[11] Non-uniform memory architecture.

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.