Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

Workloads for Niagara: Customer Deployments of Sun Niagara CMT Systems One Year On

Sponsored by: Sun Microsystems

Jean S. Bozman Matthew Eastwood

May 2007

EXECUTIVE SUMMARY

This paper explores customer adoption of the Niagara processor since its market introduction in December 2005, focusing on the range of workloads that are running today on this chip multithreaded (CMT) platform. In addition, IDC has interviewed customers who have adopted Niagara to learn about the IT requirements that Niagara addresses — and the deployment patterns found in those sites.

Sun Microsystems has innovated in the area of multicore processor technology, introducing its Niagara processors in December 2005. Each processor contains eight cores, and each of those cores supports four threads, or streams, of jobs. This means that any single Niagara processor can support up to 32 independent threads of processing, making it an efficient engine to host parallelized, multithreaded workloads. These types of workloads are typical for commercial use, such as in Web servers, Java environments, application servers (app servers), databases, and network infrastructure.

Since its launch, the Niagara-based line of systems has shown one of Sun's fastest ramps for new products. The systems have been adopted both within the Sun installed base and by new customers. Key drivers for that adoption include increased performance in more dense deployments — requiring less space within the datacenter — and smaller thermal "envelopes" due to lower levels of power dissipation. Efficiency of deployment is based on the combination of lower power and cooling requirements — and smaller space requirements within datacenter racks. Sun calls this "CoolThreads" technology due to the low power and cooling requirements of Niagara-based UltraSPARC T1 CMT-based systems, compared with earlier generations of multicore processor technology.

At the beginning, Sun took the unprecedented step of releasing the Niagara 1 processor design to the open source community via the GPL and OpenSPARC projects, allowing a new range of hardware and software developers to innovate around Sun's multithreading technology while encouraging the growth of CMT designs through a range of devices and vertical markets. In this way, Sun plans to see a rapid adoption of CMT devices that participate in the network alongside Niagara-based servers sold by Sun. In combination, large numbers of CMT systems will help to transform the network for a new generation of Web-enabled applications.

INTRODUCTION: MULTICORE AND MULTITHREADING TRENDS

The trend toward multicore processor technology is undeniable. Having started in the RISC server space in the 1990s, most RISC processors had become dual core more than five years ago — and by 2006, most of the processors for servers were built on dual-core technology. Today, a number of vendors are building or shipping quad-core processors — and the core counts are getting higher, with the aim of improving server density and performance without allowing that dense packaging to push thermal envelopes to unacceptably high levels.

Multicore technology supports a fundamental change in processor technology: Fielding single-core processors with ever-faster clock cycles inexorably led to unacceptably high costs for power and cooling of rack-optimized server systems, as well as increasingly diminishing returns for application performance.

Today, the focus is on refining that technology with the following aims:

- Reduce the thermal envelope around high-performance processor chips
- Allow some workloads to have an affinity to sets of cores, or to specific cores, for purposes of security and workload isolation and virtualization
- Support more and more threads, or streams, of work within a small, compact processor platform, which has the effect of improving overall system throughput
- Allow for wider use of multithreaded workloads, including Web-serving, video-serving application codes that have been compiled to run on multicore or multiprocessor SMP systems and workloads that already run on multithreaded operating systems
- Enable Internet workloads that are becoming increasingly parallel in nature so that individual job streams, or threads, can be directed, or assigned, to multiple cores for more efficient processing

This change in the technology supporting IT systems is having a very real business impact as well. As server density increases, more workloads can be run within a smaller space. Operational costs can be controlled by reducing the thermal footprint of IT using multicore technology while reducing the number of servers to be managed in the datacenter.

Importantly, the combination of multicore processors and multithreading allows more work to be done in smaller spaces and with lower power consumption and cooling requirements, thus preserving customers' investments in IT — and allowing them to grow capacity incrementally through the use of multiple multicore processors within one server system, one cluster, or one grid.

With so many datacenters constrained by power and space limitations, the ability to deliver the higher levels of throughput being demanded by new online computing business models is key to next-generation deployment. The higher throughput, when combined with a reduction in power and space requirements for these servers, enables organizations to deliver more services than was previously possible in traditional datacenters.

In short, Sun's Niagara-based CoolThreads servers are designed to fit into next-generation, lights-out datacenters, delivering Web-enabled services to a variety of companies, including service providers (SPs), telecommunications companies, and a range of enterprises running Web 2.0 applications.

SITUATION ANALYSIS: MARKET IMPACT OF MULTICORE, MULTITHREADED SYSTEMS

Server density is increasing, allowing more computing to be done in smaller and smaller form factors. At first, this would appear to be a contradiction in terms: How can more computing be done in smaller amounts of space? But computer science has long predicted that the trends toward increasingly powerful computing technologies, combined with computing in smaller spaces, would result in high amounts of compute horsepower being applied to a wide range of computing problems.

Why is this happening? The network is changing, and the mix of workloads flowing across the network is changing along with it. The presence of new data types, such as video, voice, and multimedia content, when combined with faster bandwidth on the network, is transforming the IT infrastructure requirements that support next-generation, Web-enabled workloads.

Video on demand (VOD), IPTV, and voice over IP (VoIP) are examples of the workloads that are supported in an interactive, multimedia-rich computing environment. The variety of workloads is increasing, bringing with it a new challenge for IT managers as well. Performance for network-centric workloads must be optimized to ensure optimal throughput. Importantly, the quality of the viewing or listening experience for end users will be affected if the processing speed for these workloads is not optimized for rapid transmission and high availability.

What can be done to improve quality of service for these kinds of workloads? Multicore processors are supporting parallel workflows, as each core takes on part of the incoming work. Multithreaded workloads can take advantage of multicore because each job or task runs independently, allowing the threads to run in separate cores.

Finally, compactness of design will help to improve overall throughput — translation: the total amount of work done on these computers. As stated earlier, computer science has long predicted that the trends toward increasingly powerful computing technologies, combined with computing in smaller spaces, would result in high amounts of compute horsepower being applied to a wide range of computing problems. The inherent parallelism in this multicore, multithreaded approach to computing improves efficiency in processing next-generation computing workloads.

But few would have predicted that high-density deployments of volume servers would result in a rapid climb in power/cooling costs. This phenomenon became clear in recent years as datacenters moved to rack-dense systems, including rack-optimized systems and bladed systems crowded into enterprise datacenter or Web-oriented colocation spaces. Now, power and cooling considerations are top of mind for IT managers — and the ability to reduce power/cooling demands, while saving on rack space, allows businesses to work more efficiently while focusing on business results.

SUN'S INTRODUCTION OF NIAGARA-BASED SERVERS

In December 2005, Sun Microsystems introduced its Niagara-based processors, each with up to eight cores and each core supporting four threads. This means that any single Niagara processor can support up to 32 independent threads of processing. At the same time, Sun introduced the T1000 and T2000 servers, which are both based on the Niagara 1 UltraSPARC T1 CMT processors.

The Sun Fire T1000 server is optimized for massively scaled computing infrastructure and, as such, packages the UltraSPARC T1 processor into one 1RU chassis typically consuming 180W of power, with pricing starting at less than $4,000. Applications include Web servers, firewalls, proxy/caches, and search farms. The Sun Fire T2000 server is packaged in a 2RU enclosure and includes greater redundancy, expandability (disk, I/O, and memory), and performance. Typical applications include virtualized and consolidated Web environments, Java application servers, enterprise application servers (e.g., ERP and CRM), and online transaction processing (OLTP) databases.

Today, Sun sells CoolThreads technology servers based on Niagara processors, including the T1000 and T2000 as well as the Netra T2000 and Netra-branded ACTA blades for telecommunications processing. All of these systems have been shown to improve performance and throughput for Solaris applications while reducing the amount of datacenter space and power/cooling needed to maintain those systems compared with more traditional server designs. When these capabilities are combined with virtualization technologies from Sun and other providers, efficiencies increase as workloads are consolidated onto Niagara-based server systems.

Sun has reported a number of customer wins across vertical market segments. In the Web 2.0 space, where datacenter efficiency is mission critical for delivery of data services, companies such as Joyent, an application development hosting site, and Fotolog, a Spanish-language social networking site and photo-sharing service, have deployed Niagara-based systems to reduce requirements for power, cooling, and rack space within the datacenter. In the financial services segment, Bank of America is using Niagara-based systems for IT infrastructure. In government, even small cities and counties are leveraging Niagara-based systems for efficient processing in smaller datacenters that typically have fewer IT staff employees and restricted IT budgets.

What does all of this mean for business? Efficient processing of Web 2.0 and content-rich multimedia workloads will speed processing for end users who are accessing these network-centric applications and other workloads, including email, databases supporting Web workloads, and a number of enterprise applications (e.g., ERP and CRM). It also saves valuable floor space in the datacenter because multiple domains of workloads are coresident on the Niagara systems, as was reported by customers who have deployed the systems.

Multicore Workloads for Niagara Servers

The "who" of the IT acquisition may ultimately be as important as the "what" when it comes to choosing Niagara-based servers or Opteron-based servers from Sun. Both are powerful volume servers, and both share a number of similar hardware components that are deployed in both types of servers. Both support lights-out and remote management features for datacenter operations. But it appears that the servers will appeal to different subsets of customers and workloads within the large Sun installed base — and outside of it as well.

One can imagine a Venn diagram showing that Sun's Niagara-based T1000 and T2000 CoolThreads servers and its Galaxy servers could support a group of workloads that run well on both platforms. This group of workloads includes support for Web serving; for Internet email and app servers such as IBM WebSphere, the Oracle app server, BEA WebLogic, and the Sun app server that is part of the Java Enterprise System (JES) of middleware; and for databases supporting OLTP workloads. There are also some workloads that would have a greater affinity for Niagara, such as audio streaming and video streaming, which can take advantage of 32 separate threads, all facing the network. That represents 32 separate pipes of data leading to the I/O channels, and ultimately to the Internet. Also, IP-based workloads, such as VoIP and the emerging IPTV, can be expected to run well on Niagara CMT processors.

Importantly, a wider range of workloads — beyond network-centric workloads alone — has been made available on Niagara-based servers, including those running on Solaris 10 and on two Linux distributions, Ubuntu and Gentoo, which are used within the Linux community for application development and network infrastructure workloads along with the familiar LAMP stack (short for the combination of Linux, Apache, MySQL, and PHP). IDC notes that Sun now offers a Solaris-based stack supporting Apache, MySQL, and PHP, which is referred to as a SAMP stack. In addition, a number of database programs are available for Niagara-based servers, including open source databases, Postgres, MySQL, and Java DB as well as other database products, such as those from Sybase, IBM, and Oracle.

Workloads: Patterns in the Data

Workloads is the word that IDC uses to describe the types of applications and tasks that are running on customer-installed servers on a daily, weekly, or monthly basis. IDC conducts annual Server Workloads research, with customer-based studies of more than 1,000 IT sites. By looking at the types of workloads that are actually running on deployed servers, IDC provides a "snapshot" of the way that servers are being used, with segmentation by operating system, by server class (volume, midrange enterprise, or high-end enterprise), and by processor type.

Figure 1 shows a workloads chart with Unix workloads for scale-up and scale-out systems. It is important to note that the full spectrum of Unix workloads is quite broad, with particular strengths in the business processing, decision support, data warehouse, business intelligence, Web infrastructure, and scientific/technical workload categories.

FIGURE 1

Worldwide Unix Server Revenue Share by Workload and Deployment Strategy, 2005

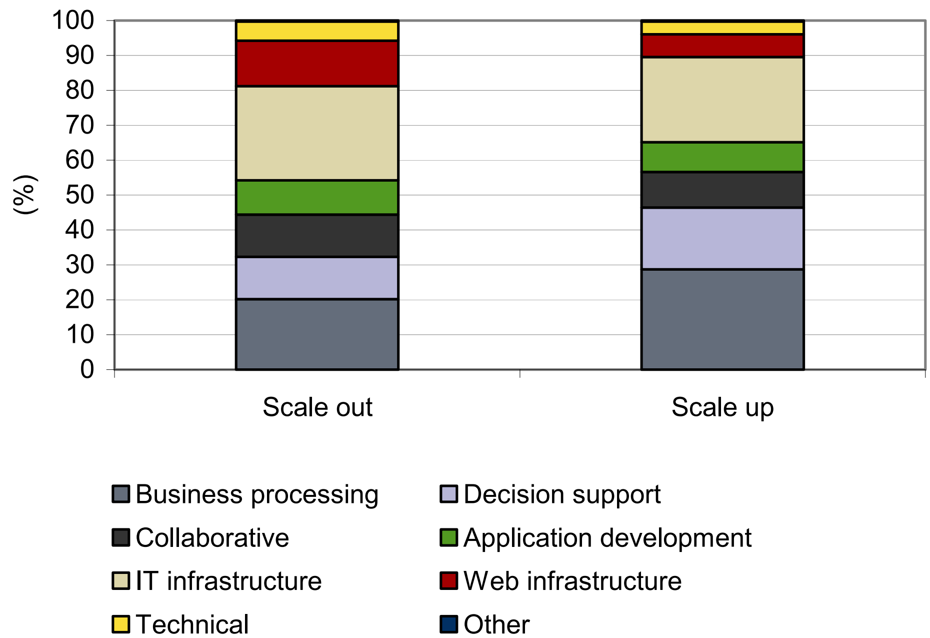

Source: IDC, 2007

A number of these workload types can benefit from running on multicore, multithreaded technology. These are discussed in the following sections.

Network-Facing Workloads

Network-facing workloads run well on multicore, multithreaded processors, given the very nature of network traffic, where packets of bits are streaming at high speeds over the Internet and private networks. Each stream of bits can be directed to one of the cores — and the throughput for this kind of workload is very high when compared with the throughput for workloads that are less "parallel," such as workloads that depend on the results of other processing in order to be updated. These network-centric workloads are in the Web infrastructure and Internet infrastructure workload categories and include not just the presentation and Web tiers, but also the application execution and database layers of the service.

Search-engine companies have a hunger for new servers — and many of them are not conventional rack-optimized servers. Rather, even the customized servers being deployed into search-engine infrastructure are optimized for speed and efficiency, with less demand for power and cooling. The workloads for search-engine companies fit into a category of workloads that IDC feels would run well on Niagara-based servers, given the servers' network-facing design, along with Niagara's enhanced support for intensive I/O operations. A search company could deploy multiple new data-based services to its end users each week, running on volume servers deployed in rack-dense configurations, where the operating temperature and power requirements are key acquisition considerations.

Service Provider Workloads

Service providers leverage a number of types of workloads, including Web infrastructure, IT infrastructure, collaborative, and business processing (e.g., ERP, CRM, OLTP, and business intelligence). SPs do so because they send data services to end users — and maintain datacenters that run the applications and data being sent over the network to those end users. Many of these SPs will maintain Web portals, on behalf of their end customers, that allow end users to access services, including Internet search.

Web 2.0 Workloads

- **Web 2.0 companies.** Companies targeting Web 2.0 workloads are looking at ways to boost throughput for data processed on the servers that are installed at their sites, or at colocation companies that support the infrastructure that runs the Web 2.0 workloads. A new range of next-generation services is coming online — all based on the protocol of the Internet, including VoIP, VOD, and IPTV. Web 2.0 companies benefit from Sun's Startup Essentials program, which wraps services and special offers around the CoolThreads T1000 and T2000 product offerings. These start-up incentives, combined with the server density and power-consumption characteristics of the Niagara-based servers, are finding new "footprints" in the Web 2.0 companies, which serve up video, email, and file-sharing services.

- **Telecommunications workloads.** Servers deployed in telecommunications companies support wireless and wireline services — and there is a need to deploy dozens, or hundreds, of rack-optimized servers within IT datacenters that are trying to keep operational costs in line. Niagara's support for telecommunications applications makes it a platform for next-generation IP-enabled workloads.

Application Development/Test Development

Application development is identified by IDC as a separate workload, one that presents its own demands in terms of the need to provide programmers with adequate compute resources for development, quality assurance, and testing of the codes that are being created. Workload isolation is also important for the application development workload because development environments can interfere with production environments on multiprocessor servers. The ability to work in parallel, within multiple domains on a single server, supports and protects app dev workloads.

Sun's Niagara Technology

When Sun Microsystems introduced its Niagara processors in late 2005, it brought to market a processor that can support up to 32 independent threads of processing. The processor contains eight cores, and each of those cores supports four threads, or streams, of jobs. Niagara is an efficient engine to host parallelized workloads, which are typical in Web, application, and OLTP database serving. The Niagara 2 (N2) processors are the second generation of this design, and they are due to ship in the second half of 2007.

Built-In Support for Solaris 10

A number of specific technologies are built into Niagara processors, including the following:

- Support for advanced virtualization and software-defined partitioning (LDoms)
- Granular control over workloads and availability of those workloads
- Optimization for the Solaris 10 operating system and its features, including the DTrace feature to identify and address system bottlenecks and Solaris Containers for application isolation
- The ZFS file system and support for advanced system management tools
- Support for Solaris 10, a fully threaded operating system with advanced thread scheduling and load balancing and CMT-optimized device drivers

Niagara processors support the SPARC V9 architecture, and are thus the CMT implementation of SPARC architecture. They are binary compatible with UltraSPARC processors and will run any applications that are deployed on UltraSPARC, preserving the investment of customers and ISVs in Solaris applications that were written to run on SPARC.

Niagara 1 and Niagara 2

Niagara 2 is the second generation of Niagara processor technology that Sun is bringing to market. It will be the basis of a range of Sun servers. Niagara 2 is a 65nm processor and will be a "server on a chip" implementation, with integration of cores, threads, backplane, memory controllers, cryptographic accelerating, networking, and I/O directly into silicon, delivering the following benefits:

- Scaling performance two times that of Niagara 1 microprocessors, which is achieved by doubling the number of threads per core and implementing two pipelines per core. This means that N2 will have 64 threads, much more than other multicore implementations.
- Doubling performance per watt, compared with first-generation Niagara systems, which is a result of scaling performance with threads rather than scaling through adding sockets or boosting frequency

- Integration of dual 10 Gigabit Ethernet (10GbE) and PCI-E controllers directly onto the N2 processor package, which reduces latency, improves I/O and network performance, and reduces the "parts count" — in turn improving reliability

- Addition of a floating-point unit and a full Visual Instruction Set (VIS) per core, which extends the benefits of Niagara-based CoolThreads servers from commercial workloads into high-performance computing (HPC) and technical workloads as well

- Integrated cryptographic acceleration in the N2 processor, delivering what Sun calls zero-cost security (Niagara 2 supports 10 common industry-standard security ciphers and is able to process them at wire speed, so there is no penalty in performance and eliminating the use and cost of plug-in crypto acceleration cards.)

Applications for Niagara

Sun estimates that a total of 1,250+ applications for Niagara can now be shown to take advantage of Niagara performance advantages. With more than 2,500 ISVs in Sun's Partner Advantage Program, representing more than 4,000 applications for Sun SPARC servers running Solaris 10 worldwide, Sun is encouraging wider application availability through its support of the Ubuntu and Gentoo distributions of Linux, along with ongoing support for the Solaris 10 operating system, which supports a wide range of Unix applications.

The CoolThreads platform has been recognized by many customers as an efficient processing platform — and it has also been recognized for its multicore support by new pricing models, reflecting this shift to a denser compute platform.

Oracle, BEA (WebLogic), IBM, SAS (Business Intelligence), and other large ISVs have changed pricing policies to provide new business models for deploying Niagara multicore processors. In December 2005, Oracle began pricing software for Niagara at a 0.25 multiplier for its Database 10g Enterprise Edition, in recognition of its high throughput and the need to offer pricing for the applications running on Niagara that is comparable to pricing being offered for other Oracle applications running on other types of processors. Since March 2007, Oracle has also redesigned its Standard edition and Standard One edition to recognize sockets as opposed to cores. This means that a fully featured Oracle database can be deployed onto Niagara with a price tag as low as $5,000.

How Niagara Supports Virtualization

When used in combination with Solaris 10, Niagara supports a unique virtualization technology that allows multiple applications to reside on a single Niagara-based server. Key to this virtualization on Niagara-based platforms is the LDom technology, which is short for "logical domain." LDoms are now available on Sun Fire T1000 and T2000 servers (and Netras driven by Niagara) with CoolThreads CMT technology.

The LDoms, which run on Sun's UltraSPARC T1 processors, support multiple Solaris 10 operating system images. Each Solaris domain "image" supports multiple Solaris Containers, each running a separate workload or application. As many as 32 LDoms can be deployed on a single Niagara processor, which is the result of mapping one LDom to one thread of workload. Sun will also have support for running Linux and BSD distributions in LDoms, side by side with Solaris instances, starting later this year.

Sun has announced two service programs aimed at improving adoption of virtualization technologies at the customer site. Solution Customer Workshops are two-day workshop sessions that assess customer requirements and generate an analysis of cost savings associated with virtualization deployments. Life Cycle Services for Virtualization are services associated with architecting, implementing, and managing virtualization technology deployments at customer sites.

Rapid Market Growth for Sun Fire Niagara CoolThreads-Based Systems

Sun has a long history of providing the server hardware that supports network-centric computing. In the case of Niagara, the idea is to support network-facing workloads that can be parallelized (e.g., video streaming, audio streaming, multimedia content sharing, Java environments, and databases in support of OLTP workloads) for rapid deployment onto this 8-core, 32-thread system. Those uses for Niagara were deployed from the beginning, just after the December 2005 launch of Niagara-based T1000 and T2000. There are also Niagara-based Netra servers for use in the telecommunications space.

IDC data shows rapid adoption of these servers, with a steep ramp in 2006, growing sevenfold in terms of unit shipments and fivefold in terms of factory revenue in 4Q06, compared with the year-ago period. The 4Q06 factory revenue represented a sequential growth rate of 26% over the 3Q06 revenue.

But streaming workloads will not be the only ones running on Niagara. Because Solaris itself is multithreaded, customers can then move existing Solaris applications, custom or ISV packaged applications, to Niagara for deployment. This can be done without changing the current code due to binary compatibility with Sun UltraSPARC systems, which Sun has maintained in the Niagara design. In fact, both UltraSPARC and Niagara are based on the same V9 version of SPARC architecture.

However, there will be several ways in which customers can optimize code to run on Niagara. One will leverage a new chipset called Neptune, introduced in March, which reroutes, or "sprays," incoming requests (coming from the Internet or the network) to parallelize them for running on Niagara CMT processors. Sun also offers software tools and compilers to help optimize codes to run even faster on Niagara than they would have without change. Sun's Cool Tools suite includes tools for application development, tuning, and deployment.

CUSTOMER SNAPSHOTS

Customers will need to factor in their IT skill sets and preferences for software deployment as they review which workloads — custom and packaged — would run on Sun T1000 and T2000 servers, based on Niagara UltraSPARC T1 CMT processors. The first generation of Niagara processors is not optimized for HPC with significant floating-point operations, or for highly scalable databases that support scalable single instances sharing large cache sizes and large memories to keep data resident on the server system while databases are maintained over time.

IDC interviewed a number of Sun customers who have deployed CoolThreads servers based on Niagara 1 technology in order to understand the workloads they are running in production today, and their experiences in using the Niagara-based technology deployed in 2006 and early 2007. To do this, Sun identified customer sites that have deployed Niagara-based servers, including the T1000 and T2000 and Netra ATCA blades based on Niagara 1 UltraSPARC T1 CMT processors. IDC analysts then used a semistructured interview guide to interview these customers. Two sites are included here.

Fotolog

Fotolog is a heavily trafficked Web site that allows Web sharing of photographs and is focused on Spanish-language and Portuguese-language commentary on those photographs. Based in New York City, the site has grown to more than 8 million members in just a few years' time and hosts 250 million photos and more than 90 million page views per day.

Fotolog is a prime example of social networking, which happens one photograph at a time, as correspondents post photos and notes that help to reduce the miles between North and South America, and between countries. There are no sharp "borders" on who may contribute because it is a worldwide Web site. But it also allows the families and correspondents sharing the photos to bridge geographical distances. There are other reasons to share: The site showcases photos taken for artistic reasons as well as personal ones. In addition, there are guidelines that support the goals of the site, prohibiting photos that show violence and nudity. The rules are simple: Post one photo a day, and share it with others who have the ability provide comments and notes about it.

At the center of the site are several dozen Sun Fire T1000 servers based on Niagara UltraSPARC T1 CMT processors. Fotolog is leveraging the multithreaded capabilities of Niagara-based servers to keep the bandwidth levels high, ensuring the quality of service to end users in Brazil, Argentina, Chile, and the United States. All of the Sun Fire T1000s at Fotolog — and there are nearly 50 of them — run Solaris 10.

Fotolog CTO Warren Habib estimates that each Niagara-based server is saving real estate in the datacenter. "On a pure performance basis, they are at minimum three times more powerful than the Sun V210s we replaced," he said. "If we didn't have 50 T1000s, we would have needed three times the number of V210s, and we would have filled five racks of servers, instead of one and a half racks." The site also has other Sun servers, including 100 Sun V210 servers, some of which run the MySQL database; several V440 servers; and eight Sun Fire V20Zs.

The Solaris 10 environment allows Fotolog developers to do programming and development within Solaris Containers, which isolate the test/development environment from the production environment. When new applications, including Java-based applications, are deployed, they move into other Containers that support production workloads. "Our development environment is identical to our production environment," Habib said, "and that makes it easier to model how everything will work." Fotolog relies on lights-out management for its Niagara-based servers, and it uses remote management to make maintenance changes, via a Web-enabled console, on a 24 x 7 basis, if needed. It can be said that the Niagara-based servers have provided a strong engine for Fotolog's core applications. "We show 72 million-plus page views per day, and we generate all those pages on the fly," Habib said, adding, "There's virtually no static content."

Joyent

Joyent, with datacenters in the San Francisco Bay Area and in San Diego, California, is a company with a new kind of business model: that of hosting application-development workloads on behalf of its end customers. Joyent's customers are typically IT organizations developing their own applications, and using Solaris 10 and its features to do so. As such, it is a company that is very much in step with Web 2.0 technologies — technologies that leverage the latest generation of Web technology and Web services as a primary component of the business model.

The Niagara-based T1000 servers running in Joyent's hosted datacenter site have proven to be efficient engines that support the company's mission-critical workloads, including the Joyent Connector capability, which in turn supports application development on behalf of end customers. "The T1000 is great for running email servers, LDAP (directory) servers, and the Postgres and MySQL database servers," said David Young CEO of Joyent. "These are network facing — really application facing — for us." The company's open source LDAP email servers leverage the multithreaded capability of the Niagara processors by being able to maintain enough connections to support a large community of end users. "The email servers are very thread aware," Young said. "It's a sort of server that has a few hundred thousand connections, and you need to be able to handle all those connections."

Some of the features of Solaris 10, running on the T1000 servers, lend themselves to troubleshooting new code, managers at Joyent reported. One is DTrace — short for Dynamic Tracing — which helps programmers and system administrators to identify any system bottlenecks that occur and allows them to pinpoint the problem and address it, removing the bottleneck condition. Another feature, the ZFS scale-out file system, allows developers to run their code against a broad and highly redundant data-file repository, speeding overall throughput because of the high degree of addressability that is provided by this extensible file system.

For many of Joyent's customers, including network equipment providers (NEPs) and other types of companies relying on Web developers, the Web-based programming world is one that supports many threads running side by side. For these customers, Young said, "Any application you write takes advantage of multithreading and parallel processing, and that's where the T1000 excels."

Importantly, the Niagara-based servers can keep pace with peaks in demand from end users. "You can burst up to 95% of the box," Young said. "As a customer, we can create more of these Solaris Containers for you, on demand, and you can scale it horizontally or vertically [by using Containers within a T1000, or by using more T1000s]." Finally, there are operational savings, with electrical costs estimated to be up to 40% less than with earlier generations of Solaris servers, saving Joyent on its utility bills. The company received some of the utility-rate savings provided by PG&E, the electrical utility in Northern California, which provides energy rebates for low-power-consuming devices.

CHALLENGES AND OPPORTUNITIES

Looking forward, the Sun Fire servers utilizing Niagara CMT CoolThreads-based systems will need to continue innovating in the marketplace, even as competitors begin their move to quad-core processors, and ever-denser processor technologies, over time. While Sun had an early start in the multicore technology space, it is by no means finished with its research in this area. Rather, Sun plans to advance its multicore processor technologies, providing more cores and threads per core in future designs and supporting more virtualization technologies along with the multicore designs. The effect of combining multicore, multithreading, and virtualization will be to provide large amounts of computing power in ever-smaller form factors, and with even higher performance per watt. This, in turn, will bring about a new generation of microsystems in Sun's server product line.

The R&D costs can be leveraged, to some extent, by working closely with other companies — in Sun's case, by working with Texas Instruments on taping out and fabricating the Niagara and Rock multicore CMT design processors. However, IDC notes that this is an active area of R&D for IBM, Intel, AMD, and others, so multicore technologies can be expected to invite close comparisons between the multicore technologies shipped — and to provide a new impetus for continuing to spend to keep ahead in terms of innovation in this technology space.

Niagara processors are highly capable, but they are not the right fit for all workloads and applications. Some workloads are not easily parallelized and run better in a more traditional SMP server, including highly scalable implementations of large databases. Others, such as intensively technical or HPC workloads, are heavily dependent on floating-point operations, and the current generation of Niagara processors (UltraSPARC T1) has not been optimized for these workloads. IDC understands that Niagara 2 processors will include a floating-point unit, and that capability has the potential to attract HPC workloads leveraging floating-point computation. However, as discussed in this paper, Niagara has been shown in its patterns of customer deployments to support many workload types.

CONCLUSION

The range of workloads that runs on Niagara has proven to be wider than was initially envisioned on first announcement and launch of the Niagara CMT processors in 2005. Customers have been able to run IT infrastructure, Web infrastructure, business processing, and collaborative workloads on Niagara — and they have done so in a range of businesses, including those in Web 2.0 start-ups, the SMB space, service providers, telecommunications companies, financial services organizations, and more traditional businesses and organizations that support Internet-enabled, network-centric workloads. They have also been able to run databases in support of Web-enabled workloads, rather than the very large corporate databases that were designed to run on scalable SMP servers.

An opportunity to extend this range of workloads exists as businesses move to modernize their aging IT infrastructure, and to Web enable the custom and ISV applications that are already installed within their corporate networks. The ability to process data with high aggregate data rates, to avoid the cost of overprovisioning rack-optimized servers by installing fewer servers, and to reduce power and cooling costs will be a factor driving some customers to acquire Niagara-based servers and follow-on products based on the Niagara 2 implementations.