

TDS

Transitional Data Services

We orchestrate IT so you can
get business *done*SM



**Sun SecureTM Application Switch – N1000 Series
Business Case**

Executive Summary

For many CIOs, three key objectives dominate the IT support agenda. They include improved organizational efficiency within stringent budget constraints, sustained or improved service delivery quality and improved business contribution, particularly in the areas of new application enablement and customer service quality.

New technology investments must concurrently support these objectives and demonstrate a compelling improvement in product state of the art. For many shops, these represent the key business case criteria for user investment.

This document discusses how Sun Secure Application Switches effectively addresses each of these requirements through a set of key benefits. We also assess the benefits provided by the N1400, the newest member of the product family.

Five key benefit categories are discussed. These include price, performance, availability, security and scalability. Related topics include the ways in which the N1400 can dramatically reduce capital and operational expenditures while improving the level of supported personalization and customer service.

An important related topic is the N1400's strong cost and performance advantages against comparable products from Cisco, Nortel and F5, consistent with other product family members. The quantitative analysis performed by Transitional Data Services, Inc. using the published price and performance figures of each of the products resulted in up to a **4.5x advantage** over these competitors in the **number of switches required to sustain specific throughput rates** with a 50%-50% mix of SSL and non-SSL traffic. Consistent with previous benchmarks, these throughput levels included 250 Megabits per second, 500 Megabits per second and 1 Gigabit per second.

Similarly, the N1400 enjoyed up to a **7x capital cost advantage** over its competitors in **sustaining similar throughput levels**. It became clear over the course of our analysis that the N1400 platform's competitive advantages are most prominent for gigabit-level traffic rates, but are also very applicable to aggregate traffic rates far less than 1 Gigabit per second.

Application Management Challenges

In order to support improved organizational efficiency, improved service delivery quality and improved business contribution objectives through new web-based application delivery, businesses must ensure the sustained presence of four key IT competencies.

The first is the ability to support mass service customization one business partner at a time. The second is the ability to both manage and deliver web-based content that both facilitates and scales the actual customization process.

The third is continuous support of secure service delivery that provides the best possible protection against compromised operations. The fourth and perhaps most important is the ability to deliver a higher level of service quality while providing substantive capital and operational cost reduction. This goes beyond doing more with less. It means doing *better* with less, especially when it comes to data center management.

Here's why. Data center budgets can often make up as much as 50 percent to 75 percent of the total IT budget, based upon data compiled over the past few years by Meta Group, Gartner Group and other recognized industry research organizations.

Given these percentages and an ongoing organizational mantra of "...doing more with less", close attention is being paid to those cost categories that make up the largest percentage of TCO.

In the case of the Data Center, the top five cost categories include support staff (often ranging between 30 percent and 35 percent of TCO), server and storage capital (between 10 percent and 15 percent), hardware maintenance (also between 10 percent and 15 percent), software maintenance (between 5 percent and 10 percent) and data center facilities (between 8 percent and 12 percent).

The need to reduce both the relative percentages and absolute values of some or all of these cost categories is one of the key forces that are driving data center consolidations. In many cases, un-managed growth of both servers and storage within individual departments and business units has often led to year-to-year TCO growth of 5 percent to 10 percent or more, significantly higher than more sustainable rates of 3 percent to 5 percent or less.

Given the need to reduce the data center TCO to more sustainable levels as well as deliver a more consistent level of IT service quality, many organizations have elected to consolidate both data center capital and staffing resources into a smaller number of centralized processing centers. This is often a significant undertaking due to a number of implementation and logistical challenges.

However, the fact is that numerous shops have achieved successful results, reducing data center TCO by anywhere from 15 percent to 25 percent or more. Obvious success factors include a clear understanding of your own TCO profile, defining realistic cost reduction targets and maintaining an unconditional focus on successful transition management.

Supporting a data center infrastructure that facilitates effective infrastructure transition, with special attention being paid to support for a scalable foundation for the addition of new business services, in addition to supporting continuous operation for existing ones, is an absolutely critical success factor.

Delivering products with superior price performance, knowing how to successfully transition and consolidate data center infrastructure as well as having substantial IT governance experience in ensuring strong IT to business alignment are all significant Sun strengths.

Sun applies sound governance principals to help ensure that its own IT operations are aligned with the business, capital and support staff resources are effectively deployed, and IT-related risks are appropriately managed. As a result, Sun has created an IT business environment that is a responsive, accountable, measurable, service-driven organization with the flexibility to adapt to ever-changing business needs.

The proof of the effectiveness of this approach is in the numbers. Based upon the results of internal measurement, IT cost reductions that lie between 12 percent and 15 percent have been the result.

In an economy in which budgets are tight and resources are scarce, having a strategic partner that is able to walk the talk relative to successfully maintaining technology leadership, having a successful project implementation track record and sustaining successful IT governance for both itself and its customers is clearly an arrow that is worth having in the operations management quiver.

Now, let's turn our attention to data center switches. We'll begin by discussing functionality and performance specifics of the N1400.

N1400 Features & Benefits

The Sun N1000 Series has been specifically developed to address the previously discussed data center price-performance, manageability and consolidation requirements.

The N1400 in particular includes faster control processors that support improved on-chip caching, faster memory, and a next generation

control chip that improves SSL connection and encrypted payload throughputs by up to 10% over previous models. In addition, application server virtualization is now an integral part of the product.

Performance Considerations

For many data center managers, however, switch price/performance is often a critical product selection criterion. When assessing switch performance, four key product metrics generally differentiate competing alternatives. These include peak data throughput, the number of concurrent connections, number of concurrent requests (or "connections") per second and the number of secure (i.e. SSL-supported) requests per second.

Peak data throughput is the raw speed in bits per second of a product's internal I/O fabric. Due to the bursty nature of most traffic workloads as well as basic physics, few, if any, products perform at or near this level of performance.

A similar argument can be made for the maximum number of concurrent connections because this metric tells the user little to nothing about how much application traffic is actively flowing over each of the individual sessions at any given point in time. In the extreme case, it can constitute an overflow point for a malicious hacker attack, such as a SYN-FLOOD.

For most users, the more relevant metrics are the number of concurrent secure and non-secure *requests per second* that the product supports. For non-secure requests, supported rates vary widely from tens up to hundreds of thousands of requests per second.

Given that encryption and decryption management are generally quite compute intensive, the supported rate of secure requests per second is typically much lower, ranging from hundreds to thousands of requests per second.

As a point of reference, consider that SSL processing can often reduce the request per second rate through a mainstream server by as much as 90%. Alleviating this specific impact is the target objective of the rapidly growing number of SSL accelerator modules and appliances.

The key to successfully understanding product performance differentiation, however, lies in a) defining exactly what makes up secure and non-secure requests in terms of protocol flows and packet contents, b) defining relevant target benchmarks in terms of request types and related flow rates and then (c) comparing apples to apples in determining the number of different vendor switch products that are required in order to sustain a given level of throughput.

What is a Request?

Fortunately, an HTTP request has a precise technical definition. This definition is generally used to characterize both the end user interaction with a web site as well as the specific TCP and HTTP protocol flows that support that interaction.

First, let's start with the user-oriented definition. In end user terms, a request begins with the entering of a URL (e.g. <http://www.yoursite.com>) to a standard browser and ends with receiving the corresponding web page content in the associated response. All underlying message types and flows constitute the protocol support.

Important and relevant considerations are product-specific performance rates that are quoted in units of connections per second metric versus those that are quoted in requests per second. By understanding the format and structure of web application transactions, one can more effectively correlate the two metrics as well as gain a more effective understanding of total traffic throughput.

The most important metric is normally *requests per second*. A request consists of a TCP connect request from the browser to the server, the forwarding of the URL to the server, the forwarding of the relevant web page content from the server to the browser and a TCP disconnect request from the browser (persistent connections that are supported in HTTP V1.1 are a separate consideration and beyond the immediate scope of this discussion).

The required TCP and HTTP connection, data transfer and disconnection overhead amount to approximately 535 bytes. The additional data required to support this round trip transaction consists of the URL string itself in addition to the corresponding web page content.

Previously, switch product throughput was quoted simply in units of *TCP connections/second*. This is obviously significantly different from the case in which the connections per second metric really refer to the number of HTTP requests per second. For this reason, the old adage of *caveat emptor* is definitely applicable.

It further behooves the user to ensure that comparable flows and metrics are being referenced when vendors are positioning their respective products. Understanding the combination of TCP and HTTP traffic that supports a quoted connection rate is absolutely crucial.

Similarly, it is important to understand the protocol subtleties of SSL requests. In one case, a significant amount of protocol handshaking needs to occur at the TCP, HTTP and SSL layers in order to support secure traffic exchanges. More specifically, the SSL flows necessary to establish a secure session and perform key exchange need to be factored into the complete set of message exchanges.

It's also important to note that an optimization can occur in which existing SSL sessions can be re-used to support new session connections. This obviates the need to re-establish the connection and re-perform key exchange set-up sequences. Obviously, the transaction performance rate of this latter case will always be far better than that of the former. Once again, it's important to ensure that what you get is what you expect.

N1400 Performance Assessment

Succinctly defining protocol message sizes, flows and handshakes in a closed form model provided the basis for an apples-to-apples comparison between the N1400 and competitive products. This model was used to characterize the ability to reliably and effectively support any mix of secure and non-secure web transactions at gigabit and multi-gigabit levels of throughput. This capability clearly differentiates the N1000 family from more conventional web switch products given that their design centers are often based upon the assumption that the vast majority of transactions are non-secure and the corresponding performance targets are on the order of hundreds of megabits per second.

In order to highlight these differences, TDS' performance analysis compared the product costs of fully configured Cisco CSS11506's (including on board SSL support provided by a separate accelerator card), F5 Big IP-6400E that supports on-board SSL acceleration and the Nortel Alteon 2424-SSL switch.

Standard vendor configuration and pricing information was used in order to perform a fair and un-biased assessment of each vendor's relative capabilities. In addition, performance results published by the Tolly Group¹ (www.tollygroup.com) and Veritest² (www.veritest.com) were used as the basis for comparison between all four switch products in order to ensure that objectively verifiable performance numbers were used as the basis for

comparison. In all cases, we assume that each vendor product can simultaneously support its quoted rates of peak SSL and non-SSL requests per second.

The CSS 11506 was measured by Veritest as supporting peak traffic rates of 4,500 new HTTP requests per second and 800 SSL requests per second per SSL accelerator card (resulting in a peak sustained rate of 3,200 SSL requests per second on a fully configured CSS 11506). A fully loaded CSS with two Gigabit Ethernet ports on its control module and 4 SSL accelerator modules, Web NS and Web NS SSH software lists for approximately \$68,375.

The F5 Big-IP 6400E switch is quoted by F5 as supporting 45,000 HTTP and 15K SSL requests per second. The associated list price of the fully configured 5100 with SSL 800 Accelerator support is approximately \$58,475.

The Nortel Alteon 2424-SSL has been measured by the Tolly Group as supporting 51,000 HTTP requests per second and the separate 2424-SSL accelerator supports approximately 2,000 SSL transactions per second. The list price of the fully configured 2424-product combination is approximately \$31,995.

Configurations of each of the four products were modeled to support varying volumes of a single HTTP request size. Our HTTP request included 3K bytes of data payload (representing the HTTP response screen) in addition to the 535 bytes worth of protocol overhead.

¹The Tolly Group, *Alteon Application Switch 2424 versus Cisco Systems CSS 11503 and F5 Networks Big-IP 5000*, January 2003

²Veritest, *Cisco Content Switching: Layer 4 and Layer 7 Performance Test*, February 2003

The 3K-byte payload was chosen due to its very representative nature as an average request size. This figure is often cited in the work of organizations such as Network Workshop and other organizations that are focused on the characterization and analysis of World Wide Web traffic.

	Sun	F5	Nortel	Cisco
250	1	1	3	2
500	1	1	5	3
1,000	2	2	9	5

Table 1—Number of Required Switches

(Source: Transitional Data Services, Inc.)

The model itself had three key variable parameters. These included the size of the HTTP request user population, the HTTP request arrival rate to the product-specific configuration and the ratio of secure to non-secure HTTP requests in the packet traffic mix.

	Sun	F5	Nortel	Cisco
250	\$25,000	\$58,475	\$95,985	\$136,750
500	\$25,000	\$58,475	\$159,975	\$205,125
1,000	\$50,000	\$116,950	\$287,955	\$341,875

Table 2—Switch Capital Cost

(Source: Transitional Data Services, Inc.)

By varying the size of the user population and HTTP request arrival rate, it was easy to determine the number of concurrent HTTP requests per second that would be required to sustain a given level of throughput. Our total throughput levels ranged from 250 megabits per second to 1 gigabit per second.

For a specific throughput level and secure to non-secure traffic ratio, our model was able to calculate the number of required switches, Fast Ethernet ports and Gigabit Ethernet ports that would be needed using the previously defined switch configuration and pricing. Specific quantitative results for the number of required switches and associated capital costs required to support 50% SSL encrypted/50% non-SSL traffic are shown below in Tables 1 and 2.

The results that are shown in Tables 1 and 2 clearly show up to a five times advantage (5x) in the number of required switches in addition to capital cost (CAP-EX) advantages that were as high as nearly seven times (7x). These results clearly show the difference between a product architecture that is optimized for secure application switching performance and more general-purpose product architectures that are very effective for lower throughput rates for largely non-secure traffic flows.

Graphical representations of the Required Number of Switches and Capital Costs versus Throughput are shown in Charts I and II respectively. These tables visually represent the relative switch number and cost differences between the four product alternatives at all major throughput levels (250 Mbps, 500Mbps and 1 Gbps).




Chart I—Throughput versus Number of Switches
 (Source: Transitional Data Services, Inc.)

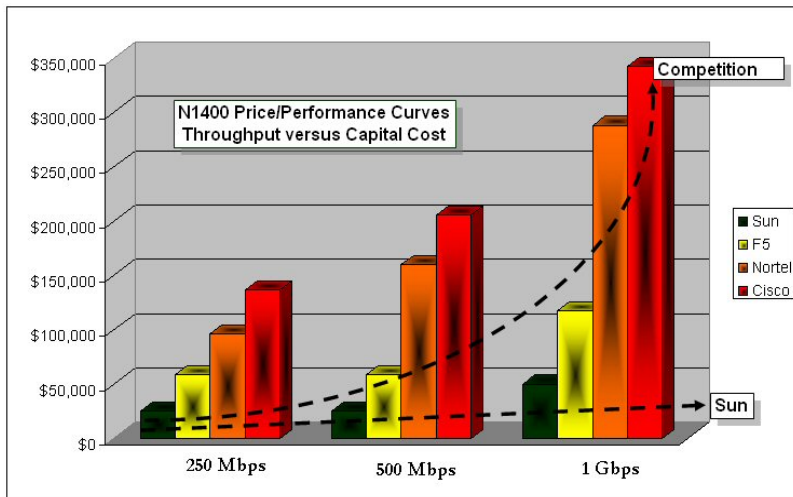


Chart II—Throughput versus Capital Cost
 (Source: Transitional Data Services, Inc.)

It is important to note, however, that Op-Ex is closely tied to both maintenance expense and outgoing cash flow, it represents a very key target for cost reduction and efficiency improvement. The other key driver of Op-EX, support staff costs, can also be positively impacted by single products that combine the functionality previously supported by separate ones, thereby reducing implementation, configuration and ongoing support complexity.

Therefore in addition to reducing aggregate maintenance cost, integrated product implementation reduces the support time required to install, configure and integrate the product support required for more sophisticated application- and content-based switching. Less support time spent on product operations means more support time that can be spent on new initiatives that can improve both the short and long term competitiveness of the business.

Customer Service Benefit

In addition to significant price-performance advantages, the combination of wire speed Application Layer processing combined with switch virtualization provides substantive customer service benefits in three key areas.

These include improved availability through logical application service partitioning, sustained performance that is largely partition-independent and an improved level of service predictability. In addition, the elimination of even one or two costly outages per year can often more than justify the capital expense, above and beyond projected price/performance benefits.

Cash Flow Benefit

Given the capital cost differences documented in the previous subsection, related operational expense (Op-Ex) reductions for product maintenance naturally fall out. This can be directly attributed to the utilization of fewer, more functional products as a replacement for legacy, less functional ones. Note that specific maintenance price percentages of product list will vary by vendor.

However, given that the great majority of maintenance price percentages lie somewhere between 15% on the low end and 20% on the high end of the capital cost of a switch, the absolute savings in annual maintenance charges will generally scale with the relative differences in capital cost.

Conclusions

Earlier in this document, we discussed the importance of support for improved organizational efficiency, service quality and business alignment in the delivery of current and future web-based applications. We also emphasized the importance of achieving these objectives while also realizing optimal value through product implementation that supports superior price/performance while also reducing Op-Ex, especially with respect to lower product maintenance cost and improved support efficiency.

In this Business case, we have shown how the N1000 Series and the N1400 in particular support these criteria using a specialized purpose-built product architecture that results in a significant price performance advantage over mainstream competition in addition to significantly reducing complexity reduction through an integrated product offering.

Through an apples-to-apples comparison of the product costs required to support representative combinations of secure and non-secure traffic at multi-megabit and gigabit rates, we've shown how the N1400 can clearly meet if not surpass these important objectives more effectively than its key competitors.

Given that the N1000 Series is the only product family that is part of a much larger application-centric product portfolio, its integration into both current and planned data centers is an option definitely worth IT decision maker consideration.

About TDS

Transitional Data Services (TDS) provides IT Orchestration services for businesses and institutions. Orchestration is an adaptive service delivery approach that significantly reduces IT management complexity while improving business alignment and the return on IT spending.

TDS combines utility-class managed services with just-in-time resource delivery in order to increase organizational responsiveness, improve service quality, and reduce operational risk. TDS' unique approach to service management enables its customers to focus on their core business and spend less time worrying about IT care and feeding.

Headquartered in Hopkinton, Mass., TDS can be reached at (877) 973-3377 or on the Web at www.transitionaldata.com.