Transitional Data Services

Sun™ Secure Application Switch—N2000 Series Business Case Study

© Transitional Data Services, Inc. 2004 All Right Reserved

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 Business Case document will describe how the N2000 product family effectively addresses each of these requirements through its support of compelling benefits.

The N2000 Data Center Switch is a purpose-built product that is specifically targeted for data center applications. Its supported combination of gigabit networking, application layer switching and secure operation (both internal and external to the product) specifically address each of the previously stated criteria.

This document defines the business case for the N2000 more broadly than previous vendor approaches. It addresses the key IT investment criteria in each of five key dimensions – Price, Performance, Availability, Security and Scalability through the use of relevant industry statistics and data as well as a design center modeling approach, this Business case will clearly present the compelling way in which the N2000 can dramatically reduce capital and, operational expenditures while improving the level of supported personalization and customer service.

Most prominent are the N2000s strong cost and performance advantages against comparable products from Cisco, Nortel and F5. The quantitative analysis performed by Transitional Data Services. Inc. using the published price and performance figures of each of the products showed a **2x to 20x advantage** over its competitors in **the number of switches required to sustain specific throughput rates with a 50%-50% mix of SSL and non-SSL traffic**. These throughput levels ranged from 250 Megabits per second up to and including 1 Gigabit per second**.**

Correspondingly, the N2000 showed up to a **4x to 10x capital cost advantage over its competitors in sustaining similar throughput levels**. It became clear over the course of our analysis that the N2000 platform's competitive advantages are most prominent for gigabit-level traffic rates, but are also very applicable to many of today's web-centric application networks.

Specific qualitative and quantitative business case details will now be discussed.

Meeting the Web-based Application Management Challenge

.

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 also providing substantive capital and operational expense 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 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.

All Right Reserved

In the case of the Data Center, the top five 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). These percentages are projected averages based upon results publicly documented by Gartner Group, Meta Group, Sun and APC.

The need to reduce the absolute values and relative percentages of some or all of these cost categories is one of the key forces that are driving data center consolidation initiatives. 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 reduction targets and maintaining an unconditional focus on successful implementation management.

In addition, supporting a data center infrastructure that facilitates effective transition, with special attention being paid to providing a solid foundation for the rollout of new business applications, in addition to supporting non-disruptive operation for existing ones, is an absolutely critical success factor. Delivering superior price performance technology, knowing how to successfully streamline and consolidate data center infrastructure and having substantial experience in successful data management and IT governance implementation are all significant Sun strengths.

According to a recent article in Inner Circle (www.sun.com/software/sunone/innercircle/) , a monthly Sun newsletter for senior executives, Sun's IT governance process has identified numerous opportunities for shaving costs and yielding millions of dollars in savings. Using a well-defined template for capturing hard-dollar benefits of IT projects, Sun can more accurately calculate how much IT projects will add to the bottom line.

Sun relies on its governance model to help ensure that IT is aligned with the business, IT resources are used properly, and IT-related risks are appropriately managed. Through governance, 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.

But the proof of the effectiveness of this model is in the numbers. According to the same Inner Circle article, Sun is achieving 12 percent to 15 percent budget impact year over year. 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 definitely an effective arrow worth having in the operations management quiver.

In the next section, we discuss functionality and performance specifics of the N2000 in more detail.

N2000 Series Overview

In any web content infrastructure, a primary objective of application switching is to more effectively deliver and personalize rapidly evolving application services through policy based forwarding on the basis of specific application identity, application content or a combination of the two. High speed switching at the application and content layers are becoming increasingly important success factors for implementing the business policy that is the foundation for highly personalized and differentiated customer service. This trend will only accelerate as Web Services technologies become more mainstream over the next few years.

 All Right Reserved

Each new product generation has been an improvement over the previous ones, yet to date not one has achieved real efficiency in effectively addressing customers' key pain points such as inadequate capacity, vulnerability to attacks, increasing demands for availability and more personalized support for key business partners. Companies doing business on the Web require technology that supports all these capabilities with significantly fewer resources in order to achieve the required levels of capital investment return.

Sun's N2000 Series Switch has been developed to specifically address these requirements. The N2000 architecture encompasses high performance network processing, a custom chipset for gigabit-scaled TCP termination and policy processing, a purpose built, mission critical real-time operating system, and customized software which provides switching, network, security, management, and virtualization application services.

The N2000 product family supports over 2 gigabits per second of either SSL or non SSL traffic, up to 250,000 new TCP connections per second, and up to 2M sustained TCP connections. Each connection delivers full TCP offload with load balancing policies ranging from more basic layer-4 request switching to full packet inspection and switching of any object of any length in the application data stream. It also supports over 12,000 new SSL connections per second and up to 480,000 concurrent SSL connections, which may be pipelined into the application and object-aware load-balancing engine.

In addition, the N2000 supports service personalization on a per customer basis. Through application layer inspection, the switch can classify arriving flows by customer. Once classified, each individual flow can then be rate-limited, policed, and if an egress port is shared, prioritized. In addition, minimum and maximum guarantees for available bandwidth as well as resources associated with TCP offload, SSL, SLB, and future object-aware switching applications are also supported.

An additional unique feature of the N2000 is its ability to support secure operations inside as well as outside the box. More specifically, the product has also been carefully designed to shield and partition data center customers from one another. To control capital costs, all users share physical resources such as enclosure, power, switch fabric, fast-path network processors, slow-path control processors, uplinks and inter-switch links. This feature is referred to as switch Virtualization. In order to ensure reliability and privacy, the Operating System assigns each customer or user group its own virtual switch (vSwitch).

Each vSwitch resides in a separate memory partition containing customer-specific forwarding information, application logic, configuration, management and security.

The primary benefits of vSwitch support are that each department, business unit and/or external customer is provided with its own secure content domain, which helps to reduce overall cost. This also ensures that each separate department, business unit and/or external customer can be supported through policies and service priorities that best address their specific business needs. In addition, vSwitch support also ensures that supported users are not unnecessarily impacted by traffic degradations, service compromises or even basic software upgrades that occur in separate domains. In addition, logical partitioning and resource sharing within a larger switch can often be far more scalable and cost effective when servicing large numbers of subscribers and/or session flows.

Key N2000 Benefits

Traffic Management

When defining cost effective infrastructure growth strategies, IT staff responsible for web application and content delivery need to carefully consider the impact of three key technologies. These technologies are Application Layer Switching, Gigabit Networking and the Secure Sockets Layer (SSL).

Let's consider the underlying rationale for each one separately.

Application Layer Switching

Like Layers 3 and 4 switching before it, Application Layer switching does not represent a radically bold departure in processing upper level protocol information. Software-based routers and specialized appliances have been processing information at and beyond Layer 4 for some time. However, Layer 4 through Layer 7 switching does represent continued progress in hardware-based switch design that allows upper layer packet processing to fully occur in tandem with wire speed Layer 2 or Layer 3 forwarding of traffic without compromising performance.

 All Right Reserved

The integration of higher level protocol intelligence within the switched LAN infrastructure greatly improves the network manager's control of the network's behavior while reducing the degree of dependence on routers, LAN Switches and earlier generation traffic management appliances, each of which supports a bounded set of functionality in addition to a product-specific network management interface.

Therefore, in order to support the four key objectives of high performance, personalization and secure processing and service quality improvement, Application Layer Switching technology is rapidly becoming the mechanism of choice for both highly flexible and intelligent content processing.

Gigabit Networking

Despite a time of increasingly tight IT spending that is only occurring for "must have" products, implementation of Gigabit Ethernet in corporate networks is continuing to grow at a very strong rate. One major factor that determines the relevance of Gigabit Ethernet to web service delivery is support on server system motherboards.

Gartner Dataquest further projects that, largely due to rapidly falling component costs and the continued need to maintain effective data center capacity, the number of shipped system motherboards that contain native Gigabit Ethernet support is roughly equal to the current shipment rate for Fast Ethernet. For most shops, data center gigabit networking is not a question of if, so much as a question of when. Gigabit Ethernet will continue the upward spiral of bandwidth consumption in enterprise applications.

Secure Sockets Layer (SSL) Support

Few argue the importance of SSL to the future of web services delivery, particularly in the support of both B2B and B2C transactions. However, use of SSL to support "client-less" service access as an alternative to more traditional IPSEC-based VPNs has rapidly grown over the past year and, for a number of good reasons, is expected to continue.

The most prominent reason is significantly reduced configuration management cost. Since SSL support is provided by every Web browser, the need to implement and configure any form of secure client support, often a mandatory pre-requisite to production implementation of IP VPNs, is vastly reduced if not totally eliminated.

In addition, secure remote access to traditional file, print and e-mail networks, the primary strength of the IPSEC VPN client, is being increasingly provided by web server gateways further solidifying the role of the browser as the universal interface engine. In addition, use of SSL enables inter-enterprise traffic to be supported more quickly since no configuration changes need to be made to the corporate firewall.

A natural outgrowth of the increasing ubiquity of SSL usage combined with the rapidly increasing implementation of Gigabit Ethernet will be the dramatic increase in the number of SSL-based transactions that need to be supported by central site web servers. In the following section, we will show how the N2000 product family has been especially optimized to meet the needs of users that do or will require highly secure yet cost effective gigabit networking support.

N2000 Price/Performance

Web Application Performance Considerations

From a performance perspective, four key product metrics differentiate web switch vendor products. These metrics are the peak throughput, number of concurrent connections, number of concurrent requests (or "connections") per second and number of secure (i.e. SSL-supported) requests per second that a given platform can support.

The peak throughput is the raw speed in bits per second of a product's internal I/O fabric. Due to both the nature of most typical workloads as well as some basic principles of 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 active traffic is 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.

 All Right Reserved

Given that encryption and decryption management are generally quite compute intensive, the supported rate of secure requests per second is quite a bit 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 available on the market.

The key to successfully understanding product performance differentiation, however, lies in a) defining exactly what makes up both secure and non-secure requests in terms of both 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 to a standard browser (e.g. http://www.yoursite.com) and ends with receiving the corresponding web page content in the associated response. All underlying message types and protocol flows are defined to support this interaction.

Important and relevant considerations are product-specific performance differences 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. We will discuss this approach in more depth in the next few sections.

The important metric is normally requests per second. From a protocol perspective, 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.

A few years ago, 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. In any case, caveat emptor is definitely applicable. It definitely 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. Once again, we'll discuss this in further detail in the upcoming sections.

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 encryption key exchange need to be accounted for in the complete set of message exchanges.

In the second case, an optimization can occur in which existing SSL sessions can be re-used to support new session connections. This obviates the need to re-do the connection and 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.

N2000-specific Application Performance Benefits

Earlier we discussed three key attributes of next generation data center switches. These were the ability to support application layer switching, gigabit networking and secure operations. In this section, we will focus on the capital cost advantages provided by the performance of the purpose-built N2000.

 All Right Reserved

We define the key performance-related business benefits of the N2000 relative to its 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 N2000 from more conventional web switch products. The design center of these products often assumes that the vast majority of transactions are non-secure and the corresponding performance targets are on the order of hundreds of megabits per second.

Performance Analysis Background & Methodology

In order to clearly illustrate 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-5100 (which also support SSL through an add-on accelerator card) and the Nortel Alteon 2424-SSL.

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 either the Tolly Group (www.tollygroup.com) or Veritest (www.veritest.com) were used as the basis for comparison between all four switch products in order to ensure that objectively verifiable performance numbers wee 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 (*ref*: Veritest, *Cisco Content Switching: Layer 4 and Layer 7 Performance Test*, February 2003) 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 accelerator cards, Web NS and Web NS SSH software lists for $73, 765.

An F5 Big-IP 5100 switch that includes an SSL accelerator card is quoted by F5 as supporting 45,000 HTTP and 800 SSL requests per second. The associated list price of the fully configured 5100 with SSL 800 Accelerator support is approximately $36,990.

The Nortel Alteon 2424-SSL has been measured by the Tolly Group (ref: The Tolly Group, *Alteon Application Switch 2424versus Cisco Systems CSS 11503 and F5 Networks Big-IP 5000,* January 2003) 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. 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.

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. 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.

 All Right Reserved

Total Throughput (Mbps)	Sun	Cisco	Nortel	F5
250	1	2	4	6
500	1	3	6	11
1,000	2	5	10	22

Table 1 – Number of Switches versus Throughput Results
(Source: Transitional Data Services. Inc.)

Total Throughput (Mbps)	Sun	Cisco	Nortel	F5
250	$37,995	$147,530	$241,980	$221,940
500	$37,995	$221,295	$ 362,970	$406,890
1,000	$75,990	$368,825	$604,950	$813,780

Table 2 – Capital Cost versus Throughput Results
(Source: Transitional Data Services. Inc.)

Graphical representations of the Required Number of Switches and Capital Costs versus Throughput are shown in Tables 3 and 4 respectively. These tables graphically show the relative cost differences between the four product alternatives at all major throughput levels (250 Mbps to 1 Gbps).

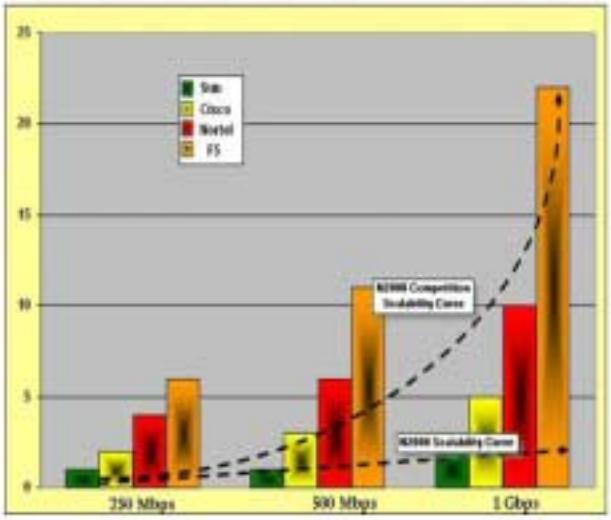

Table 3 – Total Throughput versus Number of Required Switches
(Source: Transitional Data Services. Inc.)

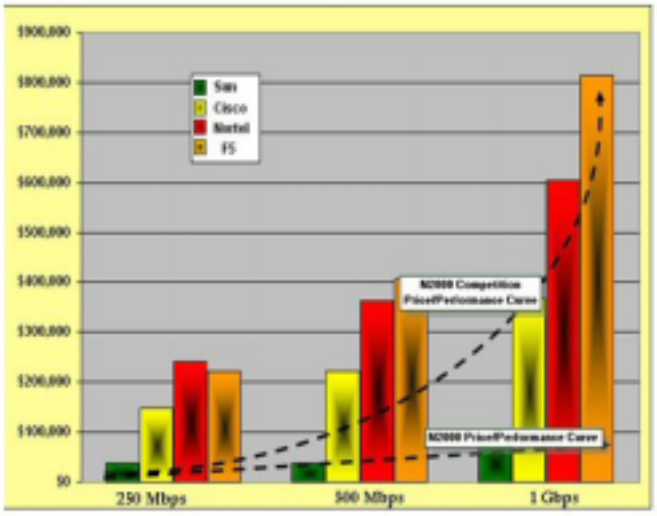

Table 4 – Total Throughput versus Required Capital Cost
(Source: Transitional Data Services. Inc.)

Reduced Capital Cash Flow Benefit Discussion

Obtaining optimum value for a given level of capital expenditure has never been more important. For many operations, IT budgets are likely to grow by no more than 3% to 5% over the coming year, especially given continuing economic uncertainty and an even stronger emphasis on affordable IT.

 All Right Reserved

Given that capital equipment expenditures often represent no more than 25% of the total IT budget, it becomes clear that dramatically new technology advances are required in order to concurrently support the earlier referenced objectives of gigabit networking, Application Layer Switching and secure operation with a close to flat capital budget allocation. In fact, it is not uncommon for IT operations to seek a capital expenditure reduction of between 25% and 45% in order to free up budgetary resources to fund new projects, making product price-performance an even more important purchase consideration.

The results that are shown in Table 2 clearly show capital cost (CAP-EX) advantages for the N2000 products that range from a minimum of three times (3x) to a high of nine times (9x), easily exceeding the aforementioned hurdle rates. These results clearly show the difference between a product architecture that is optimized for secure application switching performance and product architectures that are very effective for lower throughput rates for largely non-secure traffic flows.

The Role of Virtualization

An additional implementation challenge to consider is the following. Today, content quality, integrity and security have become must-haves as web operations become increasingly mission critical. Given that most web operations support content and service delivery on behalf of a number of separate organizations, each of which has their own unique set of requirements, the need to segregate organization-specific web services and their associated usage policy, even if for operations security reasons alone, becomes mandatory.

This is where Virtualization technology, referenced earlier in this Business Case, plays a key role. Use of VLAN technology to logically segregate content and application infrastructure was an important first step in simplifying web service operations management. However, given today's mission critical content requirements, VLAN support by itself is not sufficient. It is now extremely important to collectively support the logical and physical infrastructures, as well as related usage and security policies associated with each individual business as a single managed entity. The benefit from doing so is improved customer flexibility and self-sufficiency, as well as more secure and reliable operation.

Most of today's infrastructure products either do not support this requirement at all or they support it very inefficiently. Currently, businesses generally have to accept the level of service quality and the operations-specific customization that is provided by a switch resource that is concurrently shared with other businesses. An alternative approach is to allocate separate switches to separate businesses in order to achieve the required levels of isolation, service quality, secure operation and business-specific policy implementation. However, this is not economically practical for most operations and can easily negate the business case associated with a specific vendor technology or product family.

Meeting these requirements is what Virtualization is all about. While Virtualization is certainly necessary to address these requirements, by no means is it sufficient. High processing capacity headroom also becomes a key ingredient. The suporting rationale is that a combination of operational integrity and reliability, along with the ability to perform content domain partitioning without sacrificing the level of performance or customization required in order to optimize the customer experience, is what is really needed. Effectively meeting these requirements will be a key consideration in evolving production web service networks.

Reduced Maintenance Cash Flow Benefit Discussion

Given the capital cost differences documented in the previous subsection, related operational cost (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.

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.

 All Right Reserved

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. The integration of application and content-based forwarding together with SSL acceleration is an excellent example of just such an approach.

In addition to reducing aggregate maintenance cost, integrated product implementation reduces the support time required to install, configure and integrate the functionality 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 more directly add value to the business.

Improved Customer Service Benefit Discussion

The unique combination of wire speed Application Layer processing combined with the vSwitch capability described earlier in this document provides substantive customer service benefits in three key areas. They include improved service availability through vSwitch service partitioning, sustained performance that is largely independent of the level of individual customer personalization and an improved level of service predictability that can substantially increase the perceived quality and credibility of the service supplier. In addition, the elimination of even one or two costly outages per year can often more than justify the capital expense, above and beyond the price/performance benefits that were earlier cited.

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 N2000 Series supports 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 from a strategic vendor partner.

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 N2000 Series can clearly meet if not surpass these important objectives more effectively than its key competitors. Given that the N2000 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 considering by senior IT decision makers.

About Transitional Data Services

Transitional Data Services (TDS) provides IT Orchestration services for businesses and institutions. Orchestration is a specialized coordination IT services 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 business model enables its customers to focus on their core business and leave the complexity of IT management to experts.

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

 All Right Reserved